

Appendices

13.1 Appendix I: An Overview of Some Influential Accounts of Agent-Relativity

McNaughton and Rawling (1991, [b] 1995, [a] 1995, 1998, 2005): (x) (xS [...]). »x« stands for an agent and »S« stands for »should ensure, to the best of her abilities, and insofar as there is no conflicting duty of greater weight, that« (compare *McNaughton and Rawling* 1995b, 34). If the formulation of the relevant requirement (that is, that which the agent should ensure) contains a free agent-variable x not bound by the initial quantifier, the requirement is relative; otherwise it is neutral.

Nagel (1970): (x, A) (If R is true of A, then x has reason to promote A). To every reason corresponds a reason predicate R. The reason predicate R might or might not contain a »free agent-variable« x . If, and only if, it does contain a free agent-variable, it will be relative. A reason is relative just in case there is no corresponding neutral reason-predicate.

Portmore (2013b): »We can let » $x\varphi cEp$ « stand for » x 's φ -ing in circumstances c would ensure that p « and formulate the distinction as follows

(x)(φ)(c)(p)(if the fact that $x\varphi cEp$ constitutes a reason for x to φ in c , then this fact constitutes an agent-relative reason for x to φ in c if and only if p contains an essential reference to x , and any reason for x to φ in c that doesn't constitute an agent-relative reason for x to φ in c constitutes an agent-neutral reason for x to φ in c .)

The statement p contains an essential reference to x if and only if there is no non- x -referring statement q such that, for all x , the world in which x ensures that q is identical to the world in which x ensures that p . [...] [O]n this account, all reasons for action are facts about what agents can ensure [yet] this does not commit advocates of the account to a teleological conception of reasons, where reasons for action are all a function of what it would be good for agents to bring about. [For instance,] the fact that x 's φ -ing in c would ensure that x abides by the Categorical Imperative can constitute a reason to [sic!] x to φ in c even if x 's abiding by the Categorical Imperative would produce bad consequences.« (*Portmore* 2013b, 165 and 166.)

Ridge (2011): For every reason there is some default principle such as

(P) For all possible agents (x), all possible actions (φ) and all facts (F) If F is a fact to the effect that x 's φ -ing would promote pleasure and no other feature of the situation explains why F is not a reason to φ then F is a reason to φ .

Then: »The default principle corresponding to a given reason will either include a *non-trivial* free-agent variable in the statement of the reason or not. If it does then the reason is agent-relative; otherwise it is agent-neutral. The idea is that the use of a free-agent variable to indicate that the action is one available to the agent for whom the fact is a reason is trivial in the sense that it must be included in the statement of any reason whatsoever. So the reasons associated with the default principle (P) are agent-neutral, as the only use of the free-agent variable (x) is the trivial one that indicates that φ is a possible action of p 's.« (Ridge 2011, 20f.) Note: on Ridge's account all reasons are facts about something's being true of some agent's action.¹⁵⁹

Rønnow-Rasmussen (2009): »If P states a reason for x to φ , then: P states an agent-relative reason for x if and only if P contains an essential reference to x , otherwise P states an agent-neutral reason for x to φ . Moreover, P contains an essential reference to x if and only if P is not logically equivalent with any other statement Q that does not refer to x .« (229f.)

Skorupski (1995): $P\varphi -> x$ has reason to φ . P is the reason predicate, φ is an action (or inaction) open to the agent, x is the agent. Again, if P contains a free occurrence of x (and if it cannot be expressed by a neutral predicate without change in content), then the corresponding reason is relative.

13.2 Appendix II: On Mark Schroeder's Objection to Accounting for Deontic Restrictions against the Backdrop of a Fitting-Attitude Analysis of Agent-Relative Value

Mark Schroeder (2007) attributes the following claims to a fitting-attitude analysis of relative and neutral value (or, as he calls it, good-relative-to and good simpliciter):

a is better simpliciter than b if, and only if, it is fitting for everyone to prefer a over b .
 a is better-relative-to x than b if, and only if, it is fitting for x to prefer a over b .

159 About his account's compatibility with particularism Ridge says: »(P) is compatible with the particularist's holistic conception of reasons. For in those cases in which the status of a fact about pleasure as a reason is defeated by sadism (e.g.) the »no other feature of the situation explains why F is not a reason ...« clause is not satisfied.« (Ridge 2011, 19.)

Suppose a defender of agent-relative consequentialism wanted to account for a deontic restriction against killing the innocent against the backdrop of such a fitting-attitude analysis of value. Then, as Schroeder points out, it should hold good that, first, Franz's murdering is better simpliciter than his not murdering if his murdering prevents Hans and Jens from murdering. For there would be less murders overall. Second, however, it should also hold good that Franz's not murdering is better-relative-to Franz than his murdering. For only if his not murdering is better-relative-to Franz than his murdering, we can make sense of why Franz should not murder rather than murder. However, these two claims seem to be incompatible if one assumes a fitting-attitude analysis of value of the kind offered by Schroeder.

For according to such an analysis, the claim that Franz's murdering is better simpliciter than his not murdering will entail the claim that it is fitting for *everyone* to prefer Franz's murdering over his not murdering. If it is fitting for *everyone* to prefer Franz's murdering over his not murdering, it will also be fitting for *Franz* to prefer his murdering over his not murdering. The problem is that this last statement directly contradicts what we get when we unpack the second of the two claims above, according to which Franz's not murdering is better-relative-to Franz than Franz's murdering. On the proposed fitting-attitude analysis of value, this last claim entails the claim that it is fitting for Franz to prefer his not murdering over his murdering. Thus, it is fitting for Franz to prefer both his murdering (a claim that follows from the judgment about what is better simpliciter) and his not murdering (a claim that follows from the judgment about what is better-relative-to Franz). And that seems to be flat out inconsistent.

Both Smith (2003, 2009) and Suikkanen (2009) have offered – as I find, convincing – replies to this objection. Since both replies rest on the same basic idea and since I find Smith's portrayal of that idea clearer, I will concentrate on his proposal. The crucial point of Smith's reply is that we need to distinguish between what he calls value-making features and evaluative facts. Evaluative facts are facts about goodness. Evaluative facts are relative just in case they are facts about what is valuable relative to whom; they are neutral just in case they are facts about what is valuable in an unqualified sense. That is, neutral evaluative facts are facts about Moorean goodness; relative evaluative facts are facts about goodness conceived of as a relational property. For some reason or other – which need not interest us here – Smith thinks that we should conceive of *all* evaluative facts as being relative.

Value-making features, by contrast, are those ›things‹ – properties, I suppose – that make it the case that something is valuable. Value-making features are relative just in case their specification contains a reference to the agent relative to whom they (their instantiation) is valuable; neutral otherwise. An example might help to illustrate. If Franz's being happy makes it the case that a state of affairs in which Franz is happy is good-relative-to Franz, then the relevant value-making feature (Franz's being happy) is relative. If, by contrast, someone's being happy makes it the case that a state of affairs

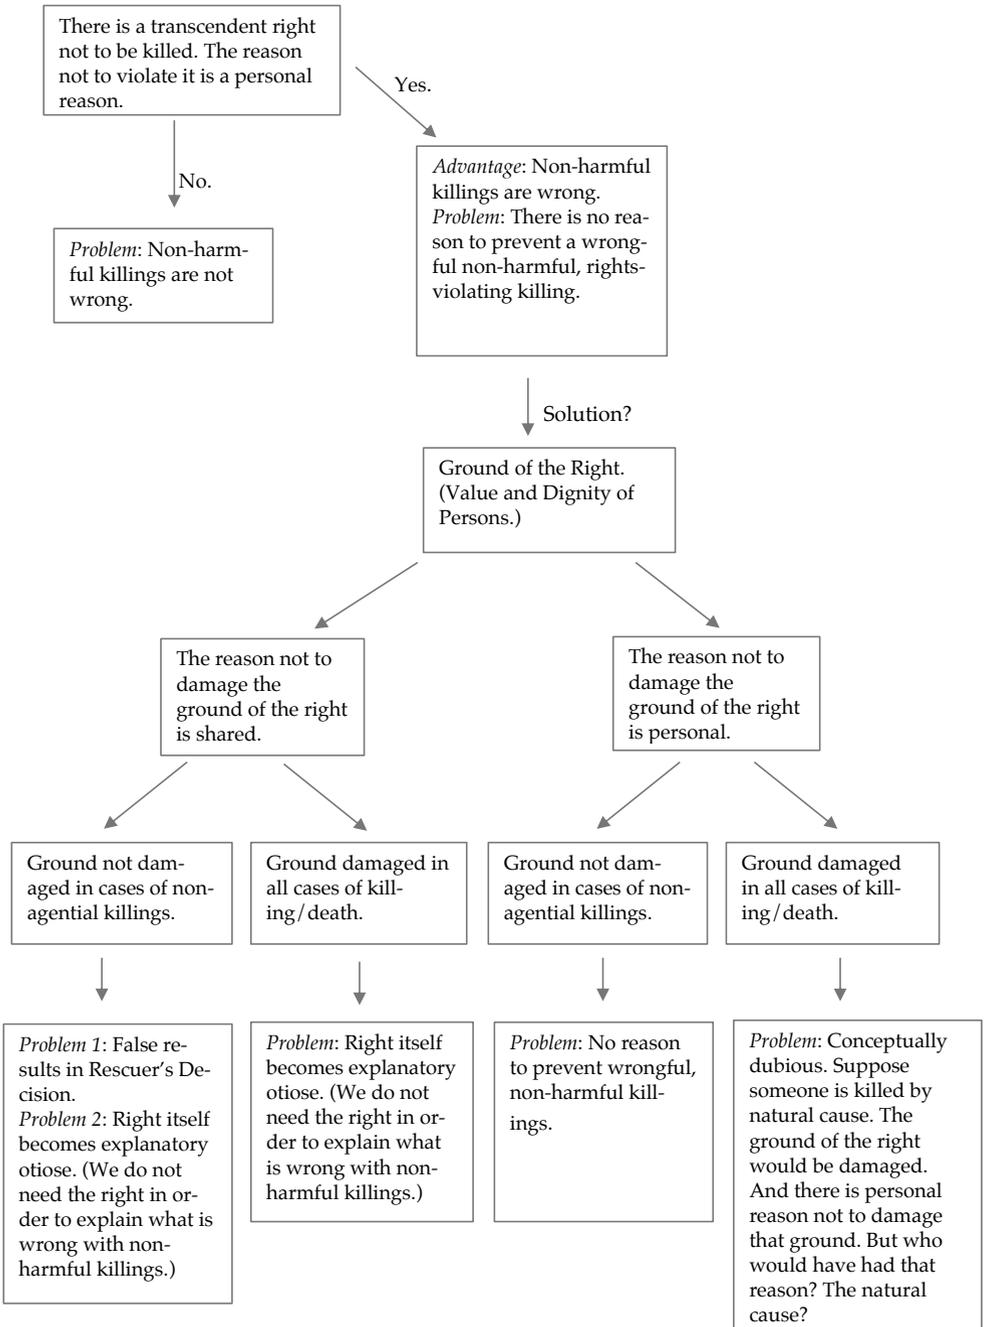
in which someone is happy is good-relative-to Franz, then the relevant value-making feature (someone's being happy) is neutral. Since Franz's being happy only contributes to what is good-relative-to Franz, whereas someone's being happy contributes to what is good-relative-to anyone, this example also illustrates that relative value-making features are relevant only for evaluative facts about what is good-relative-to the agent to whom there is reference in the specification of the value-making feature.

How does all this bear on Schroeder's objection to accounting for deontic restrictions against the backdrop of a teleological framework that avails itself of a fitting-attitude analysis of value? Well, the problem of Schroeder's objection is that it locates the distinction between relative and neutral value at the level of evaluative facts, whereas it should be located at the level of value-making features. For the sake of the argument, let us assume, as Smith does, that all evaluative facts are relative. Thus, there will be facts about what is better-relative-to, but no facts about what is good *simpliciter*. In that case, a fitting-attitude analysis of value can still run the distinction between neutral and relative value as follows:

a is better-relative-to x than b if, and only if, it is fitting for x to prefer a over b. a is *neutrally* better-relative-to x than b if, and only if, the specification of the feature that makes it fitting for x to prefer a over b does not contain reference to x. a is *relatively* better-relative-to x than b if, and only if, the specification of the feature that makes it fitting for x to prefer a over b does contain reference to x.

With this analysis in mind, let's return to the case of Franz, Hans, and Jens. According to that analysis, if Franz's murdering is agent-neutrally better(-relative-to Franz) than Franz's not murdering, this means that it is fitting for Franz – and everyone else – to prefer there being fewer murders rather than more. And, according to that analysis, if Franz's not murdering is agent-relatively better(relative-to Franz) than his murdering, this means that it is fitting for Franz to prefer that he, Franz, not murders rather than murders. Thus, it is fitting for Franz to prefer both there being fewer murders rather than more and his not murdering over his murdering. While these preferences pull into different directions, so to speak, when it comes to the question of whether Franz should murder in order to prevent Hans and Jens from murdering, there is nothing inconsistent about its being fitting to have both of them. (Remember that the upshot of Schroeder's analysis was that it is fitting to have inconsistent preferences.)

13.3 Appendix III: Full List of Possibilities regarding the Ground of the Right Not to Be Killed and the Reasons Not to Damage that Ground



Transcendent right not to be killed: There is a right that is different from and transcends the fact that it is impermissible to kill someone due to the harm this causes and due its damaging the value and dignity of the person killed. If there is such a right, there will be a reason not to kill someone that differs both from the reason not to kill someone due to the harm this causes and from the reason not to kill someone due the killing's damaging the value and dignity of the person killed.

Ground of the right not to be killed: The ground of the right is what justifies the right. Assumption: ground of the right not to be killed = the value and dignity of persons.

13.4 Appendix IV: Full List of Possibilities of What is Valuable about Being Inviolable

Suppose that what is valuable about inviolability might be our being publicly recognised as inviolable (by legal or social norms or the behaviour of others), our believing to be inviolable, our *de facto* being inviolable, or any combination of these factors. Let REC stand for »being publicly recognised as«, BEL for »believing to be inviolable«, and BEI for »being inviolable«. Then the following table gives us all possible combinations of what might be valuable about inviolability.

	REC	BEL	BEI	What is valuable is ...
1.	REC	BEL	BEI	... being recognised as, believing to be, and <i>de facto</i> being inviolable.
2.	REC	BEL	¬ BEI	... being recognised as and believing to be inviolable.
3.	REC	¬ BEL	BEI	... being recognised as and <i>de facto</i> being inviolable.
4.	REC	¬ REC	¬ BEI	... being publicly recognised as inviolable.
5.	¬ REC	REC	BEI	... believing to be and <i>de facto</i> being inviolable.
6.	¬ REC	REC	¬ BEI	... believing to be inviolable.
7.	¬ REC	¬ REC	BEI	... <i>de facto</i> being inviolable.
8.	¬ REC	¬ REC	¬ BEI	–

Line (8.) obviously is a non-starter. If nothing about being inviolable is valuable, then there is no sense in which being inviolable is valuable.

Lines (2.), (4.), and (6.) all fall in one category in the following sense: according to the interpretations of IA that are represented by these lines, the value of inviolability can be realised without our *de facto* being inviolable. Thus, the »argument from inviolability« would not be an argument from inviolability but an argument from the value of certain designs of social or legal institutions or an argument from the value of having certain beliefs. Since I assume that the argument from inviolability is meant to allude to the status of *being* inviolable, we can put these interpretations to one side.

Remain (1.), (3.), (5.), and (7.).

13.5 Appendix V: On Doing and Allowing Optimal and Non-Optimal Harm

Here is how I want to understand the notions of doing harm and allowing harm to happen, respectively:

Doing Harm: If you would have acted differently than you in fact did and would have done nothing instead and some person would, in that case, been better off than she now in fact is, then, by acting as you did, you have harmed that person.

Allowing Harm to Happen: If there was an option, besides doing nothing, that you could have seized instead of the option you in fact seized and if, in case you would have seized the other option, some person would have been better off than she now in fact is, then, by acting as you did, you have allowed harm to happen to that person.

Further, here is what I mean by optimal harm and non-optimal harm, respectively:

Optimal and Non-Optimal Harm: Particular harm the occurrence of which reduces global harm is optimal harm. Particular harm the occurrence of which increases global harm is non-optimal harm.

Against this backdrop, here are some examples of doing optimal harm, doing non-optimal harm, allowing optimal harm to happen, and allowing non-optimal harm to happen:

Example 1: A and B each need €5 to realise a project of theirs. C needs €10 to realise a project of hers. A and B do not have €5, respectively. C has €10 Euro. The realisation of all projects would have the same value (for the agents as well

as impersonally conceived). If A or B takes away the €10 from C and distributes them between A and B, this would be an instance of *doing optimal harm*.

Example 2: As above. Except that, this time, A and B each have €5 and C does not have €10. If C steals A's and B's money and uses it for his project, this would be an instance of *doing non-optimal harm*.

Example 3: Same as example 1. Except that, this time, there is a bystander D. By not intervening, she *allows optimal harm to happen*.

Example 4: Same as example 2. Except that, this time, there is a bystander D. By not intervening, she *allows non-optimal harm to happen*.

Both a constraint on optimal and a constraint on non-optimal harming allow agents to let both optimal and non-optimal harm happen. This is so because both constraints are only constraints on *doing harm*. They do not prohibit allowing harm to happen in any way.

The constraints differ, however, in their stances towards *doing harm*. A constraint on optimal harming always prohibits doing harm, irrespective of whether the harm is optimal or non-optimal. That is, a constraint on optimal harming would not permit what is done in example 1 and in example 2.

By contrast, a constraint on non-optimal harming prohibits doing harm only if doing harm is non-optimal. As a consequence, a constraint on non-optimal harming permits what is done in example 1 (A's or B's stealing C's money) but not what is done in example 2 (C's stealing A's and B's money).