

Contractualism and Deontic Restrictions*

Jeffrey Brand-Ballard

Dilemmas with the following structure play a central role in normative ethics: an agent must decide between doing harm to one person and allowing harm to befall several others. That is, he must decide whether to perform a *minimizing violation*.¹ Moral theories diverge regarding when, if ever, minimizing violations are permitted. Direct act consequentialism, for example, often permits and sometimes requires them: the agent may maim another person whenever doing so is expected to promote the good (e.g., by minimizing total maimings). Most of us, however, believe it is wrong for the agent to maim in at least some such cases. This is to say that our ordinary morality includes certain *deontic restrictions* (“restrictions”).²

Defenders of restrictions (“restrictionists”) disagree among themselves about the nature of restrictions and the proper way to incorporate them within moral theory. More disturbingly, several philosophers have

* This article expands on papers presented to the Minnesota Philosophical Society (St. Joseph, 2001), the Pacific Division of the American Philosophical Association (San Francisco, 2003), and the International Society for Utilitarian Studies (Lisbon, 2003). I thank my Minnesota commentator, Michelle Mason, and the other attendees, especially Norman Dahl. My APA commentator, Paul Hurley, supplied challenging feedback and kindness, for which I am grateful. I also learned from my ISUS commentators, Tim Mulgan and Roger Crisp. I am indebted to Brad Hooker for inviting me to join his ISUS discussion panel, for arranging the panel, and for offering great encouragement. I benefited from early conversations with Nomy Arpaly, George Sher, Baruch Brody, and Russ Shafer-Landau. Peter Vallentyne wrote detailed comments which clarified my thinking. Thanks to an anonymous reviewer and three associate editors at *Ethics*, David DeGrazia, and Laura Brand-Ballard for helpful feedback and enthusiasm. Derek Parfit kindly shared some of his forthcoming work as this article was going to press. John Deigh offered encouragement and suggestions. I alone am responsible for remaining defects and errors.

1. I take the term from Kasper Lippert-Rasmussen, “Moral Status and the Impermissibility of Minimizing Violations,” *Philosophy & Public Affairs* 25 (1996): 333–50.

2. I use the term ‘deontic restriction’ rather than the more common ‘agent-centered restriction’ or ‘agent-relative restriction’ so as to remain neutral on the controversial question of whether deontic restrictions are, indeed, agent centered. See Eric Mack, “Deontic Restrictions Are Not Agent-Relative Restrictions,” *Social Philosophy and Policy* 15 (1998): 61–83.

Ethics 114 (January 2004): 269–300

© 2004 by The University of Chicago. All rights reserved. 0014-1704/2004/11402-0003\$10.00

worried that restrictions are deeply paradoxical—"intuitions in search of a foundation," in Samuel Scheffler's phrase.³ The list of usual skeptics includes, in addition to Scheffler, Shelly Kagan, Derek Parfit, Jonathan Bennett, and David Cummiskey ("restriction skeptics").⁴

The skeptics' challenge has met many responses.⁵ I shall concentrate on a line of response growing out of the accusation that restrictions only appear paradoxical to the skeptics because the latter are caught in the grip of certain undefended metaethical prejudices.⁶ In short, the story goes, the skeptics make the mistake of identifying the impersonal standpoint with the ethical standpoint. This way of responding to the skeptics has inspired some interesting recent attempts to defend re-

3. Samuel Scheffler, *The Rejection of Consequentialism*, 2d ed. (New York: Oxford University Press, 1994), p. 80. The worry originates in Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic, 1974), pp. 28–35, though Nozick ultimately favors restrictions.

4. Scheffler, *The Rejection of Consequentialism*, chap. 4; Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1989), chaps. 5, 9; Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), chap. 4; Jonathan Bennett, *The Act Itself* (Oxford: Oxford University Press, 1995), chap. 10; and David Cummiskey, *Kantian Consequentialism* (Oxford: Oxford University Press, 1996), and "Kantian Consequentialism," *Ethics* 100 (1990): 586–615. Kasper Lippert-Rasmussen has shown that different skeptics characterize the paradoxical nature of restrictions in subtly different ways ("In What Way Are Constraints Paradoxical?" *Utilitas* 11 [1999]: 49–70). Roughly, the skeptics reason as follows. The deliberate maiming of an innocent person is an intrinsically bad event. If maiming an innocent is so bad as to be forbidden, then how can a practice be worse if it keeps these bad events to a minimum? If we understand rightness in terms of bringing about good states of affairs and/or avoiding bad ones, then it is hard to see how restrictions could be anything but irrational. This is the alleged paradox.

5. See, e.g., F. M. Kamm, *Morality, Mortality* (Oxford: Oxford University Press, 1996), vol. 2; Stephen L. Darwall, "Agent-Centered Restrictions from the Inside Out," *Philosophical Studies* 50 (1986): 291–319; Christine M. Korsgaard, "The Reasons We Can Share: An Attack on the Distinction between Agent-Relative and Agent-Neutral Values," *Social Philosophy and Policy* 10 (1993): 24–51; Thomas Nagel, *Equality and Partiality* (New York: Oxford University Press, 1991), chap. 13; Alec Walen, "Doing, Allowing, and Disabling: Some Principles Governing Deontological Restrictions," *Philosophical Studies* 80 (1995): 183–215. Some argue that, however paradoxical restrictions appear to the restriction skeptics, they enjoy such overwhelming intuitive support that we should incorporate them directly into our moral theory as axioms and give up trying to justify them in terms of other considerations. If we cannot derive strong restrictions from more basic norms, they conclude, this just suggests that attempts to derive ordinary morality from more basic norms are misguided. This inference counsels, perhaps, a return to pluralistic intuitionism. Jonathan Dancy seems to adopt this strategy in *Moral Reasons* (Oxford: Blackwell, 1993). See also Philippa Foot, "Utilitarianism and the Virtues," *Mind* 94 (1985): 196–209; Robert M. Stewart, "Agent-Relativity, Reason, and Value," *Monist* 76 (1993): 66–80; and W. D. Ross, *The Right and the Good* (Oxford: Clarendon, 1930).

6. Restrictionists who make this accusation include Paul Hurley, "Agent-Centered Restrictions: Clearing the Air of Paradox," *Ethics* 108 (1997): 120–46; Rahul Kumar, "Defending the Moral Moderate: Contractualism and Common Sense," *Philosophy & Public Affairs* 28 (2000): 275–309; and Michael Ridge, "Saving Scanlon: Contractualism and Agent-Relativity," *Journal of Political Philosophy* 9 (2001): 472–81.

restrictions. Several of these attempts employ T. M. Scanlon's hypothetical contractualism.⁷ Until the appearance of these efforts, contractualism played no more than an implicit role in debates concerning restrictions. However, Scanlon insinuates that restrictions might be justifiable in contractualist terms.⁸ Regrettably, he declines to elaborate, but Michael Ridge has reiterated Scanlon's suggestion,⁹ and others, notably Paul Hurley and Rahul Kumar, have tried to make good on it.¹⁰

These proposals demand special attention because many have found contractualism a powerful unifying framework within which to understand various other features of ordinary morality. It would be a result of great philosophical significance to show exactly how contractualism supports restrictions, as well. It would bring together two major strands of deontological ethics—one metaethical and one normative—moving deontology toward an important consolidation.

More than a decade ago, a restriction skeptic, Shelly Kagan, speculated that “no plausible version of the contract approach would be adequate for a defense of ordinary morality.”¹¹ But Kagan did not claim to have demonstrated this. This article examines the extent to which contractualism can, in fact, support restrictions—a central feature of ordinary morality. My minor thesis is that contractualism is a patient-focused theory. By this I mean that contractualism does not ultimately assign special normative weight to the features of agents qua agents. My major thesis is that, insofar as contractualism is patient focused in this sense, it cannot support restrictions which are strong enough to answer the skeptics.

Section I explains the restrictionists' response to the skeptics. Section II shows how restrictionists and their critics often talk past one

7. See T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998), and “Contractualism and Utilitarianism,” in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press, 1982), pp. 103–28. Hurley, Ridge, and Kumar all invoke Scanlon.

8. Scanlon, *What We Owe to Each Other*, pp. 81–86, 106–7.

9. Ridge, p. 477.

10. Hurley; Kumar, “Defending the Moral Moderate.” A defense of restrictions in terms of Rawlsian contractualism appears in Alexander Rosenberg, “Contractarianism and the ‘Trolley’ Problem,” chap. 8 of *Darwinism in Philosophy, Social Science and Policy* (Cambridge: Cambridge University Press, 2000). More pessimistic assessments appear in articles by William N. Nelson and Arthur Applbaum. Nelson examines the possibility of supporting the Doctrine of Double Effect, which embodies a deontic restriction, in terms of a version of contractualism (not there identified as Scanlonian). Nelson concludes, as I do, that the prospects are discouraging (“Conceptions of Morality and the Doctrine of Double Effect,” *Journal of Medicine and Philosophy* 16 [1991]: 545–64). Applbaum reaches similar conclusions, also without discussing Scanlon by name, in “Are Violations of Rights Ever Right?” *Ethics* 108 (1998): 340–66.

11. Kagan, *The Limits of Morality*, p. 46, n. 37.

another. I conclude that contractualism can, indeed, support certain types of restrictions. Unfortunately, these restrictions turn out to be relatively weak in certain respects which I shall make precise. Weak restrictions prove highly sensitive to empirical contingencies, losing their force under adverse empirical conditions. Hence, they are unlikely to satisfy many deontologists. What is more important, in this dialectical context, I find that these restrictions are decidedly weaker than those attacked by the restriction skeptics. None of the skeptics denies that such weak restrictions are defensible, for example, in terms of suitably sophisticated forms of consequentialism. The skeptics reject only strong restrictions, those that cannot be defended in terms of any form of consequentialism.¹² Section III argues that contractualism cannot support stronger restrictions than consequentialism can.

Some restrictionists have recognized that strong restrictions are the actual and exclusive targets of the skeptics. These restrictionists have, appropriately, sought to defend strong restrictions against the challenge. Perhaps they will yet succeed. I offer an interim conclusion: Scanlonian contractualism, despite its many virtues, offers nothing distinctive to restrictionists, even in its most developed form. To this extent, Kagan's pessimism about the capacity of contractualism to support ordinary morality gains some credibility.¹³ Restrictionists might want to explore other conceptions of impartiality or agent-centered approaches such as virtue ethics.

I. IMPERSONALITY AND IMPARTIALITY

Restrictionists such as Hurley, Ridge, and Kumar consider contractualism a promising basis for restrictions because they accept a certain critique of the restriction skeptics. That critique suggests that restrictions only appear paradoxical because the skeptics try to force restrictions into an impersonal conception of value. Hurley's discussion is representative. He observes that the skeptics recognize only two evaluative standpoints, the personal and the impersonal. The impersonal, agent-neutral standpoint is appropriate for the global evaluation of states of affairs. The personal standpoint is the agent's own standpoint, from which she determines how well things are going for her. The skeptics assume that the rightness of an action is a function of the global evaluation of the states of affairs that the action is expected to bring about. Since the impersonal standpoint is appropriate for the evaluation of states of affairs, and the rightness of actions depends on the evaluation of states of affairs, according to the skeptics, they conclude that the

12. Later, I will refer to these as 'nonminimizing restrictions'.

13. Assuming that ordinary morality really does favor strong restrictions, not just weak ones.

impersonal standpoint must also be appropriate for the moral evaluation of actions. The skeptics simply identify morality with the impersonal standpoint.¹⁴ They assume, without argument, that the evaluation of states of affairs is the only relevant consideration in evaluating actions as right or wrong.¹⁵ The skeptics' excuse for this assumption is that no one has offered any alternative to the impersonal standpoint as the moral standpoint. Hurley concludes that what "purports to be an argument that there is no rationale for agent-centered restrictions is . . . an argument that no such rationale can be reconciled to a thoroughly impersonal conception of value."¹⁶

Hurley thinks this is a mistake because there is, in fact, a third standpoint that is neither personal nor impersonal. We can visualize it as located somewhere in between these poles. This third standpoint is that of impartiality.¹⁷ In taking up the impartial standpoint, one does not abstract completely from one's own personal perspective. But neither does one aim to take the "perspectiveless perspective" associated with the impersonal standpoint.¹⁸ The impartial standpoint does not eliminate each of our individual perspectives, it merely requires us to eliminate bias.¹⁹ Impartiality does not necessarily generate only reasons to promote good states of affairs, either directly or indirectly. Therefore, even though restrictions sometimes forbid promoting the good, they are not necessarily paradoxical from an impartial standpoint. Hurley suggests that previous restrictionists have been, at least implicitly, assuming an impartialist conception of value.

Hurley, along with Kumar and Ridge, has done us a service by clarifying the potential connection between restrictions and impartiality.

14. Hurley, p. 122.

15. *Ibid.*, pp. 130–31. Ridge accuses three reviewers of Scanlon's book of making this same mistake; see esp. p. 481. The reviews criticized are Colin McGinn, "Reasons and Unreasons," *New Republic*, May 24, 1999, p. 35; Philip Pettit, "Doing Unto Others," *Times Literary Supplement*, June 25, 1999, pp. 7–8; and Simon Blackburn, "Am I Right?" *New York Times*, February 21, 1999.

16. Hurley, pp. 143–44.

17. *Ibid.*, pp. 132–37. He identifies Scanlon as a representative impartialist. Other writers refer to impartiality as an "interpersonal" or "intersubjective" standpoint and model it in somewhat different ways than Scanlon. See Stephen L. Darwall, *Impartial Reason* (Ithaca, N.Y.: Cornell University Press, 1983); Korsgaard; Brian Barry, *Justice as Impartiality* (New York: Oxford University Press, 1995); Elizabeth Anderson, *Value in Ethics and Economics* (Cambridge, Mass.: Harvard University Press, 1993). In this article I limit myself to Scanlonian contractualism. I shall not consider whether restrictions are defensible in terms of other conceptions of impartiality. 'Contractualism' refers to Scanlon's theory, in its various incarnations.

18. See Samuel Freeman, "Contractualism, Moral Motivation, and Practical Reason," *Journal of Philosophy* 88 (1991): 281–303, p. 303.

19. Hurley, p. 145. Hurley never defines 'bias', but I read him to mean something like "unjustified preference in favor of oneself or one's favorites."

I agree with these authors that the restriction skeptics have typically assessed restrictions in isolation from broader metaethical commitments. The skeptics overlook the possibility of defending restrictions in impartial terms, and the paradox of restrictions does, indeed, derive from a feature of the skeptics' conception of value.

Notice, however, that tracing the paradox to the skeptics' conception of value and offering a competing conception does not suffice to resolve the paradox. In addition, restrictionists must show that their conception of value does not share with the skeptics' conception the very features which generate the paradox. Otherwise, their attack on the skeptics' conception of value fails in its purpose.

So we have three tasks. First, we must analyze more closely the skeptics' conception of value. Then we must isolate which features are responsible for the paradox. Finally, we must determine whether the restrictionists' conception of value shares those features.

The conception of value which restrictionists attribute to the skeptics appears derived from utilitarianism. We can split this conception into four distinct theoretical commitments. The first of these is a *value-maximizing conception of rationality*, according to which rational action is action that maximizes value.²⁰ Second, the conception includes an *instrumental conception of morality*, according to which the purpose of morality is to bring something about.²¹

Third, and most important for my purposes, we have *patient focus*. A patient-focused theory can take as its objects actions, rules, states of affairs, and so on. A patient-focused theory evaluates its object exclusively in terms of the object's actual and/or hypothetical implications for individuals qua patients—in their capacities as individuals who endure and are acted upon. Patient-focused theories contrast with agent-focused ones. The latter evaluate their objects at least partly in terms of the features of individuals qua agents—in their capacities as performers of actions or promoters of states of affairs.

The final commitment attributed to the skeptics is *welfarism*. According to welfarism the only facts that carry intrinsic, objective value are those that concern the welfare of individuals. Welfarism constitutes a specific version of patient focus, according to which the only implications that matter for an individual qua patient, are those concerning her welfare.²² Objective value theories, by contrast, accord intrinsic, objective value (instead of, or in addition) to factors which are not components of anyone's welfare.

20. See Samuel Scheffler, "Agent-Centered Restrictions, Rationality, and the Virtues," in *The Rejection of Consequentialism*, 2d ed., pp. 142–43.

21. See, e.g., Korsgaard, p. 24.

22. Familiar varieties of welfarism include hedonism and desire-satisfaction theories.

Utilitarians accept all four of these commitments: value-maximizing rationality, instrumental morality, patient focus, and welfarism. Nonetheless, they are distinct commitments. Contractualists, and other impartialists, sometimes fail to differentiate them.²³ I am not sure why. Perhaps, at a subconscious level, they assimilate consequentialism to utilitarianism (consciously, of course, they know better).

Whatever the explanation, failure to draw these distinctions leads contractualists astray. They predict that the paradox of restrictions will not arise under their alternative conception of value—impartiality. At first, this prediction enjoys some credibility. Contractualism does, indeed, abandon the value-maximizing conception of rationality and the instrumental conception of morality. Scanlon has recently abandoned welfarism as well, in his important book, *What We Owe to Each Other*.

Nevertheless, I shall argue that the paradox of restrictions does not depend on the value-maximizing conception of rationality, the instrumental conception of morality, or welfarism. Though these commitments may be sufficient for the paradox to arise, they are not necessary. Rather, the paradox derives from the remaining commitment, patient focus. Contractualism is, by design, as patient focused as consequentialism. Contractualism asks whether any patient affected by a proposed rule would have an agent-relative reason to reject it. In some cases these reasons derive from the extent to which the rule would diminish the patient's welfare. In other cases the reasons do not concern welfare. For example, the rule treats the patient unfairly, or sends a message authorizing her mistreatment, or permits her to be misused. In any case, for contractualists, whether or not someone has a reason to reject a rule is always a function of her agent-relative reasons. Moreover, the relevant agent-relative reasons are not agent-relative reasons of the individual qua agent—in her capacity as one who acts on rules.²⁴ Rather, they are agent-relative reasons of everyone qua patient—in one's capacity as one who does not act on rules but is otherwise affected by them.²⁵ This is the sense in which Scanlon's theory is patient focused.

These observations connect my analysis to the distinction, now common in the literature, between agent-focused and patient-focused ra-

23. Hurley, e.g., simply lumps them together as the skeptics' dubious "conception of value" (p. 122).

24. This interpretation stands at the center of Ridge's recent defense of Scanlon. And Ridge reports that Scanlon himself has confirmed this interpretation in correspondence (p. 475, n. 10).

25. Kamm makes a similar point when she notes that "contractualist theories (like Thomas Scanlon's) . . . claim that to raise an objection to a proposed policy an individual must do so in virtue of the effects of the policy on him alone." F. M. Kamm, "Précis of *Morality, Mortality*, vol. 1: Death and Whom to Save from It," *Philosophy and Phenomenological Research* 58 (1998): 939–45, p. 942.

rationales for restrictions.²⁶ Agent-focused rationales support restrictions in terms of the individual's responsibility for her own conduct, as an agent to whom restrictions are addressed. Patient-focused rationales support them in terms of the patient's standpoint—how restrictions might affect her or how she might respond to them, not as one to whom restrictions are addressed. A defense of restrictions is either agent-focused, patient-focused, or neither. David McNaughton and Piers Rawling have argued that patient-focused theories cannot support strong restrictions.²⁷ If they are correct that patient-focused theories cannot support strong restrictions, and if I am correct that contractualism is a patient-focused theory, then we have additional grounds for doubting the viability of contractualist defenses of strong restrictions.

II. HARM-MINIMIZING AND POSITIVE-HARM-MINIMIZING RESTRICTIONS

I now introduce the first of several scenarios for discussion:

In scenario 1, Kat, Liz, and Mel are partially trapped on the train tracks, unable to extricate themselves. Each member of this innocent trio will lose an arm under the oncoming train unless Ben, a bystander, cuts off the arm of another bystander, Jennifer, and

26. See David McNaughton and Piers Rawling, "Deontology and Agency," *Monist* 76 (1993): 81–100. Deontologists are currently divided over whether to understand restrictions in agent-focused or patient-focused terms. Defenses of agent-focused rationales include Darwall, "Agent-Centered Restrictions from the Inside Out"; Eric Mack, "Moral Individualism: Agent-Relativity and Deontic Restraints," *Social Philosophy and Policy* 7 (1989): 81–111; Anderson, p. 29; Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986); Foot, "Utilitarianism and the Virtues." Others defend patient-focused rationales, e.g., Richard Brook, "Agency and Morality," *Journal of Philosophy* 88 (1991): 190–212; F. M. Kamm, *Morality, Mortality* (Oxford: Oxford University Press, 1993), vol. 1, and "Harming Some to Save Others," *Philosophical Studies* 57 (1989): 227–60. In many cases, a restriction dictates the same results, regardless of how we justify it. However, it is now well-documented that cases exist in which a restriction based on a patient-centered rationale delivers the opposite result from a restriction based on an agent-centered rationale. And even when the results are the same, a patient-focused rationale gives different reasons for the result than does an agent-focused rationale. See McNaughton and Rawling, "Deontology and Agency."

27. McNaughton and Rawling, "Deontology and Agency." These authors do not distinguish clearly between strong and weak restrictions. They present their conclusion as though it applies to all restrictions. I think they should have concluded, more cautiously, that patient-focused theories can support, at most, what I shall call positive-harm-minimizing restrictions—the same restrictions that can be justified in terms of an axiom that always permits optimizing. I think this is what they meant.

wedges it into the track mechanism in order to switch the train to another track. Ben can stop the train in no other way.²⁸

A. Harm-Minimizing Restrictions

A rule that forbids an agent (e.g., Ben) to maim (injure, kill, deceive) an innocent, unconsenting bystander (e.g., Jennifer) is a *restriction rule*. These rules can play different roles within our moral theory. We can distinguish between moral axioms and corollaries. An *axiom* is an objective, ultimate, nonderivative criterion of right. As such, it is appropriate for use by detached observers and evaluators. A *corollary* is a norm justified in terms of a more basic norm, such as an axiom or another corollary.²⁹ Distinguishing between axioms and corollaries makes room for two-level (and *n*-level) moral theories, such as rule consequentialism, “sophisticated” consequentialism, and contractualism itself.³⁰

Which axioms support restriction rules as corollaries? Consider axioms that always permit promoting the impersonal good. These are *optimization-permitting* axioms.³¹ Although optimization-permitting axioms never directly forbid promoting the good, many writers have explicitly defended restriction rules as corollaries of optimization-permitting axioms.³²

28. This scenario adapts the Trolley Problem from Philippa Foot, “The Problem of Abortion and the Doctrine of Double Effect,” in *Virtues and Vices* (Berkeley: University of California Press, 1978). I use maiming rather than killing in order to avoid philosophical distractions associated with the unique disvalue of death. Ben faces a choice between “intentionally doing” and “foreseeably allowing” harm. If he maims Jennifer, he avoidably violates not one but two deontic restrictions: one against intentionally harming and one against doing harm. This makes the scenario a good paradigm. If restrictions make sense anywhere, they should make sense here.

29. Some corollaries resemble axioms in that they, too, are objective criteria of right. Other corollaries are subjective decision procedures, appropriate for use by agents in the process of deliberation. I shall ignore the distinction between objective and subjective corollaries. See, e.g., Henry Sidgwick, *The Methods of Ethics*, 7th ed. (Indianapolis: Hackett, 1981), p. 413; R. E. Bales, “Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?” *American Philosophical Quarterly* 8 (1971): 257–65.

30. Peter Railton coins the term *sophisticated consequentialism* in “Alienation, Consequentialism, and the Demands of Morality,” *Philosophy & Public Affairs* 13 (1984): 134–71, p. 153.

31. Optimization-permitting axioms include consequentialism, which always requires promoting the good, and “hybrids,” which always allow, but do not require, promoting the good (i.e., they incorporate “prerogatives” or “permissions” not to optimize). See Scheffler, *The Rejection of Consequentialism*.

32. See, e.g., Philip Pettit, “The Consequentialist Can Recognize Rights,” *Philosophical Quarterly* 38 (1988): 42–55; Christopher McMahon, “The Paradox of Deontology,” *Philosophy & Public Affairs* 20 (1991): 350–77; G. F. Schueler, “Consequences and Agent-Centred Restrictions,” *Metaphilosophy* 20 (1989): 77–83; and Cummiskey, *Kantian Consequentialism*, pp. 12–14, 151, n. 1.

Optimization-permitting axioms can support only the least restrictive type of restrictions, which I call *harm-minimizing restrictions*. To understand the latter, we need the familiar, though problematic, distinction between doing and allowing (e.g., killing vs. letting die). Happily, my purposes do not require me to analyze this distinction. I shall simply write that the agent positively harms the patient if he does the harm to her rather than merely allowing harm to occur (e.g., the agent throws an infant into a well). The agent negatively harms the patient if he merely allows harm to occur (e.g., the agent finds an unknown infant drowning in the well and does not rescue her when he could do so).³³ I shall use the term *harm*, simpliciter, to refer, collectively, to both positive and negative harm.

Harm-minimizing restrictions forbid positive harming, even if positive harming would minimize harm in that case, unless a widely accepted practice of positive harming in such cases would minimize harm, simpliciter, in the long run.³⁴ An optimization-permitting axiom can support harm-minimizing-restriction rules, given certain pervasive empirical contingencies. These contingencies—familiar from the rule-utilitarian literature—include bounded rationality, informational limitations, slippery slope factors, and so on.³⁵

B. Positive-Harm-Minimizing Restrictions

Most deontologists, however, will complain that harm-minimizing restrictions are too weak. Some object that harm-minimizing restrictions fail to capture the special urgency that ordinary morality associates with its prohibitions on killing, maiming, and so on.³⁶ They want something stronger. They also want restrictions that are less sensitive, or not sensitive at all, to empirical contingencies. At the very least, deontologists want *positive-harm-minimizing restrictions*.³⁷ These forbid positive harming unless a widely accepted practice of positive harming in such cases would minimize harm (and other bad events) in the long run, with extra weight assigned to positive harm to reflect its additional intrinsic badness.

Might contractualism impose something stronger than harm-min-

33. See, e.g., Bennett, chap. 4; Warren Quinn, "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing," *Philosophical Review* 98 (1989): 287–312.

34. On my definition, even a harm-minimizing restriction prohibits positive harm in many instances in which positively harming maximizes total well-being, or minimizes harm, in that case.

35. Michael Byron, "Satisficing and Optimality," *Ethics* 109 (1998): 67–93; Philip Pettit, "Consequentialism and Respect," *Ethics* 100 (1989): 116–26; John Broome, *Weighing Goods: Equality, Uncertainty, and Time* (Oxford: Blackwell, 1991).

36. See, e.g., Ross, p. 35.

37. See, e.g., F. M. Kamm, "Harming Some to Save Others," pp. 251–52; Quinn, pp. 311–12; Anderson, pp. 73–90.

imizing restrictions, such as positive-harm-minimizing restrictions? Within a contractualist framework, to say that a restriction is valid is to say that someone could reasonably reject a rule that did not include the restriction in favor of a restriction rule as a basis for the general regulation of behavior.³⁸ Scanlon now tells us that the reasonableness of rejecting a rule is a function not just of its effects on individuals' welfare but also of independent moral considerations such as responsibility, fairness, rights, and entitlements.³⁹ This "generic reasons" model gives the contractualist much needed flexibility, compared to Scanlon's earlier, "welfarist" model.⁴⁰ Scanlon can now avail himself of important factors identified by standard deontology.⁴¹ For example, he can invoke considerations concerning how harms come about in addition to considerations of welfare.⁴²

The generic reasons model presupposes nonwelfarist value theory. This enables the contractualist to support positive-harm-minimizing restrictions. Consider, in particular, *deontic value theories*, objective value theories that accord intrinsic, objective value to factors such as what intentions the agent has, whether harms are positively or negatively connected to human agency, whether harms occur as means to an agent's end or merely as side-effects, and so on. According to some deontic value theories, being positively maimed, as such, gives the victim a stronger reason for rejection than being negatively maimed in the same way would give her.⁴³ In scenario 1, Jennifer has a reason to reject a rule that imposes no harm-minimizing restrictions, while Kat, Liz, and Mel have a reason to reject a positive-harm-minimizing-restriction rule. However, since Jennifer is threatened with positive harm, while Kat, Liz, and Mel are threatened with only negative harm, Jennifer has a stronger reason to reject her disfavored rule than they have to reject theirs. This is to say that Jennifer can reasonably reject a rule that imposes no harm-minimizing-restrictions in favor of a positive-harm-minimizing-restric-

38. In Scanlon's words, "an act, A, is wrong if and only if any principle that permitted A could reasonably be rejected by people moved to find principles for the general regulation of behavior that others, similarly motivated, could not reasonably reject (or, equivalently, if and only if A would be disallowed by any principle that such people could not reasonably reject)" (*What We Owe to Each Other*, p. 4).

39. *Ibid.*, pp. 204–5.

40. *Ibid.*, p. 217; see also Derek Parfit, "Equality and Priority," *Ratio* 10 (1997): 202–21. But see Sophia Reibetanz, "Contractualism and Aggregation," *Ethics* 108 (1998): 296–311, p. 311, n. 18, for concerns about the new model.

41. I take the phrase "standard deontology" from F. M. Kamm, "Owing, Justifying, and Rejecting," *Mind* 111 (2002): 323–54, p. 340.

42. *Ibid.*, p. 339.

43. Although not all positive harm is worse than all negative harm.

tion rule. Hence, the generic reasons model can support positive-harm-minimizing restrictions.⁴⁴

However, it is important to recognize that all the work, not just some of it, is being done by the shift to deontic value theory, which the generic reasons model presupposes.⁴⁵ This becomes evident when we consider optimization-permitting axioms that incorporate the same value theory. Such an optimization-permitting axiom can forbid a positively harmful act even if a practice of positive harming in such cases would minimize total harm in the long run because, according to deontic value theory, a single positive harm is intrinsically worse than (some) finite sets of negative harms numbering more than one. Therefore, an optimization-permitting axiom plus deontic value theory can support positive-harm-minimizing restrictions, too.⁴⁶ To this extent, contractualism has no evident advantage over an optimization-permitting axiom with respect to the defense of positive-harm-minimizing restrictions.⁴⁷

III. NONMINIMIZING RESTRICTIONS: THE REAL PUZZLE

It should not surprise us to find that optimization-permitting axioms can support even positive-harm-minimizing restrictions. Restriction skeptics, including Scheffler, Kagan, and Bennett, acknowledge the possibility of defending certain restrictions as corollaries of optimization-permitting axioms.⁴⁸ They decline to specify the strength and scope of

44. We can even imagine a *lexical* deontic value theory that classifies positive harms of a certain severity as both intrinsically and *incommensurably* worse than negative harms of the same severity. Lexical deontic value theory gives *absolute* priority to avoiding positive harm. A contractualism featuring lexical deontic value theory could support absolute harm-minimizing restrictions and absolute positive-harm-minimizing restrictions. Scheffler mentions “pluralist lexical consequentialism” in *The Rejection of Consequentialism*, p. 27.

45. By attributing “deontic value theory” to contractualists I do not suggest that deontic considerations form part of their theory of *prudential* value, although this is how a consequentialist would incorporate deontic considerations. My point is just that the generic reasons model can support restrictions stronger than harm-minimizing restrictions only by allowing that deontic factors supply us with reasons to reject rules.

46. Combined with lexical deontic value theory, an optimization-permitting axiom could even support absolute positive-harm-minimizing restrictions, because that value theory regards a single positive harm as intrinsically worse than any number of negative harms.

47. Kumar recognizes that a contractualist defense of restrictions should offer some advantage over, e.g., indirect consequentialist defenses (“Defending the Moral Moderate,” pp. 276–77).

48. Hurley notes this in “Agent-Centered Restrictions,” p. 121. Scheffler repeatedly emphasizes that he is not attacking indirect consequentialist defenses of restrictions. He justifies his choice by appeal to his belief that “traditional deontological views . . . maintain . . . that . . . there is a rationale for agent-centered restrictions which is independent of even broadly consequentialist considerations.” He also states that the “prohibitions imposed by rule consequentialism and related ‘two-level’ views do not count as agent-centered” in his sense (“Deontology and the Agent: A Reply to Jonathan Bennett” in *The*

the restrictions that could be so defended, but I expect they would include positive-harm-minimizing restrictions in this category.

So which type of restrictions do the skeptics actually challenge? The only outstanding puzzle concerns the most restrictive type: nonminimizing restrictions. These forbid positive harming in at least one case in which a widely accepted practice of positive harming in such cases would minimize positive harm in the long run.⁴⁹ It does appear that optimization-permitting axioms cannot support these, no matter how creative we get with our value theory. To illustrate, let us add a new character to our scenario:

Scenario 2 is identical to scenario 1, except that Sue is operating the train and can stop it harmlessly, if she wishes. She will use the train to maim Kat, Liz, and Mel unless Ben maims Jennifer, just as severely.

In scenario 2, Kat, Liz, and Mel are threatened with positive harm, just as is Jennifer. Imagine that, given prevailing social conditions, a practice of positively harming in cases such as scenario 2 is expected to minimize positive harms in the long run.

The transformation of scenario 1 into 2 exemplifies a highly general skeptical strategy, which I call *equalization*.⁵⁰ In whatever way the restrictionist characterizes Ben's action in maiming Jennifer, the skeptic tries to use that same characterization for what Sue threatens to do to Kat, Liz, and Mel (modifying the hypothetical facts, as necessary). The interesting question is: for any given generic reason which Jennifer has to reject the minimizing-violation-permitting rule, can the skeptic equalize the hypothetical facts, and the description thereof, such that someone else (Kat, Liz, or Mel) has a generic reason, of at least equal weight, to reject the nonminimizing-restriction rule?⁵¹ The restrictionist must argue that Jennifer has a generic reason to reject the minimizing-violation-permitting rule and that no one else has a generic reason, of at

Rejection of Consequentialism, 2d ed., p. 155, n. 2). Elsewhere he entertains as "not . . . implausible" the idea of derivative restrictions ("Prerogatives without Restrictions," in *The Rejection of Consequentialism*, p. 181). Similarly, Kagan recognizes that a "two-level" moral theory, such as indirect consequentialism, could support a highly qualified restriction against harming others, though he emphasizes that such restrictions would not apply whenever the agent was certain that more good would result from ignoring them (*The Limits of Morality*, pp. 32–39). See also Bennett, pp. 22–26.

49. Hereafter I omit the qualifier 'widely accepted'. On my definitions, even a minimizing restriction prohibits positive harm in many instances in which positively harming maximizes total well-being, or minimizes harm, in that case.

50. See related discussion of the "consequentialist vacuum cleaner" in David McNaughton and Piers Rawling, "Agent-Relativity and the Doing-Happening Distinction," *Philosophical Studies* 63 (1991): 167–85, p. 167.

51. 'Minimizing-violation-permitting rule' is just a less awkward substitute for 'nonminimizing-restriction rule'.

least equal weight, to reject the opposite rule. If the restrictionist can show this, then Jennifer can reasonably reject the minimizing-violation-permitting rule. Otherwise, the paradox of nonminimizing restrictions recurs.

Nonminimizing restrictions forbid positive harming even in cases such as 2. An optimization-permitting axiom, however, has no rationale for categorically forbidding such positive harms. An optimization-permitting axiom still supports corollaries that minimize harm, even taking into account that positive harm is intrinsically worse than negative. Such an axiom cannot support a rule that categorically forbids one positively to harm when a practice of doing so would minimize positive harms in the long run (and have no other bad effects).⁵² This is why I conclude that no optimization-permitting axiom supports nonminimizing restrictions. However, it remains to be seen whether contractualism can support nonminimizing restrictions—the only type of restrictions that optimization-permitting axioms cannot support—any more successfully than optimization-permitting axioms can.

Here it is important to be generous in our treatment of contractualism. In their defense of restrictions, contractualists may make liberal use of other moral concepts, including the resources of standard deontology. I shall purposely resist the temptation to charge contractualism with circularity the instant it makes use of independent deontological concepts.⁵³ We can think of contractualism as a unifying framework within which, it is claimed, the resources of standard deontology can be more effectively deployed in the defense of restrictions.⁵⁴ A contractualist defense of restrictions only stands convicted of circularity if we find that all the work is being done by the independent moral concepts to which the contractualist appeals. This is what we discovered with respect to harm-minimizing restrictions and positive-harm-minimizing restrictions. It remains to be seen whether this is the case with respect to nonminimizing restrictions.

Does contractualism deliver on its promise to defend nonminimizing restrictions? Or do the characteristic rigidities of contractualism

52. In some cases, otherwise similar to 2, inflicting positive harm causes, in addition to its desirable effect of preventing further positive harms, some undesirable indirect effects in the form of negative harm. That negative harm could be so great as to outweigh the preventative benefits of the positive harm, in which case even an optimization-permitting axiom would forbid inflicting that harm. In 2, however, Ben's maiming of Jen has no such remote bad effects.

53. Scanlon discusses circularity at *What We Owe to Each Other*, pp. 194–95. See also Brad Hooker, "Scanlon's Contractualism, the Spare Wheel Objection, and Aggregation," in *Scanlon's Contractualism: Readings and Responses*, ed. Matt Matravers (London: Frank Cass Publishers, 2003).

54. Scanlon, *What We Owe to Each Other*, pp. 213–18.

actually hinder the deployment of otherwise useful deontological resources? In the rest of this section, I shall argue as follows. Some familiar deontological factors, such as fairness, intention versus foresight, and the need to treat persons as ends in themselves, prove useless to any defense of nonminimizing restrictions. I argue that these concepts become no more useful when incorporated into a contractualist framework.

Other concepts hold out more promise, in their own right, as resources for defending nonminimizing restrictions. These include the following ideas: (1) the individual has a special claim on her own body; (2) no one has the right to commandeer someone else's body for his own purposes; and (3) the individual should be regarded as enjoying an inviolable status. However, I shall argue that these more useful concepts all function poorly within a contractualist framework. Ironically, contractualism neutralizes these deontological resources rather than catalyzing them.

The reason, to anticipate, is that a theory cannot support nonminimizing restrictions if it defines the agent's moral reasons ultimately in terms of agent-relative reasons possessed by anyone *qua* patient. To support nonminimizing restrictions, a theory must define the agent's moral reasons partly in terms of his properties *qua* agent—such as his particular mental states or basic actions. Contractualism does not.

A. *Fairness*

Consider Scanlon's sensible treatment of rescue crises.⁵⁵ Scanlon admits that contractualists need to explain why a rescuer faced with two groups of accident victims, one larger than the other, is usually obligated (and is certainly permitted) to save the larger if he can save only one group. The rescuer may not, for example, simply flip a coin to decide which to save. Scanlon observes that failing to give priority to the larger group would amount to making the decision just as though the groups had been of equal size. Making the decision on that basis effectively treats some members of the larger group as though their presence made no difference whatsoever. Those members could object that their interests have been ignored, hence they could reasonably reject a principle that fails to give the larger group priority.⁵⁶

This plausible endorsement of aggregation in rescue contexts, however, puts additional pressure on contractualists to justify nonminimizing restrictions. We have assumed that members of the larger group of rescue victims can reasonably object that their interests have not been taken into account by a rule that authorizes a decision by coin flip. If this is so, then surely the victims of a nonminimizing-restriction rule

55. *Ibid.*, pp. 229–41.

56. *Ibid.*, p. 232.

can reasonably object that their interests have not been taken into account by a rule that positively guarantees they will suffer harm. So the argument Scanlon uses to defend aggregation in the rescue context threatens to justify aggregation, as well, in other contexts in which we must harm some to save others.⁵⁷

B. Intention

Many standard deontologists analyze scenarios such as 2 in terms of the allegedly nonderivative moral significance of the distinction between intention and foresight. According to the Doctrine of Double Effect, maiming Jennifer is worse than refusing to do so because Ben intends Jennifer's injury if he maims her, but merely foresees the injuries to Kat, Liz, and Mel if he refuses. Perhaps contractualists can make use of this distinction, too. They could claim that Jennifer has a stronger reason to reject a minimizing-violation-permitting rule than Kat has to reject the opposite rule because Jennifer is threatened with intended harm, while Kat is only threatened with foreseen harm.

This reasoning rests on an illusory asymmetry. Although Ben does not intend harm to Kat, Liz, or Mel, these three will nevertheless endure intended harm if Ben refuses to maim Jennifer. Sue does not merely foresee but positively intends her actions to harm them. The harm to Kat looks unintended only if we limit our focus to Ben's agency. The fact that Ben intends to harm Jennifer does give Jennifer a reason to reject a minimizing-violation-permitting rule. But the fact that Sue intends to harm Kat gives Kat just as strong a reason to reject the contrary rule. This looks like one of the cases Thomas Nagel warns us about, where every available rule seriously disadvantages someone, yet no one can reasonably reject any of the alternatives.⁵⁸

It is interesting to note, moreover, that Scanlon himself has recently disavowed the Doctrine of Double Effect.⁵⁹ Scanlon would agree that Ben is wrong to maim Jennifer (and not wrong to refuse), but he denies that Ben's intentions are basic right-making factors. He offers a different explanation (see Sec. III, E below).

57. On Scanlon's latent acceptance of aggregation, see Michael Otsuka, "Scanlon and the Claims of the Many versus the One," *Analysis* 60 (2000): 288–93; and the responses of Rahul Kumar, "Contractualism on Saving the Many," *Analysis* 61 (2001): 165–70; and Iwao Hirose, "Saving the Greater Number without Combining Claims," *Analysis* 61 (2001): 341–42. See also Alasdair Norcross, "Contractualism and Aggregation," *Social Theory and Practice* 28 (2002): 303–14; Reibetanz.

58. Nagel, *Equality and Partiality*, p. 49.

59. T. M. Scanlon, "Intention and Permissibility," *Proceedings of the Aristotelian Society* 74, suppl. (2000): 301–17.

C. Persons as Ends-in-Themselves

Next consider how a contractualist might try to adapt Kant's Formula of Humanity. She could suggest, plausibly, that the fact that a certain rule would result in someone's not being treated as an end-in-herself (and/or her being used as a means only) gives her a reason to reject the rule. When Ben maims Jennifer, he uses her as a means to an end—the end of saving Kat, Liz, and Mel. Arguably, he also fails to treat her as an end-in-herself. Perhaps these facts give Jennifer a generic reason to reject the minimizing-violation-permitting rule. By contrast, it would seem that a restriction rule does not lead Kat, Liz, or Mel to be treated as means only. The minimizing-violation-permitting rule permits someone to be treated as less than an end-in-oneself to be used as a means only. The restriction rule, by contrast, permits no one to be so used. The restrictionist could insist on account of this difference that the former rule is reasonably rejected in favor of the latter. This looks like a basis for a restriction rule.

This conclusion is too quick. So far I have not specified the purpose that Sue would serve by injuring Kat, Liz, and Mel. But I can stipulate that Sue also intends to treat them simply as means to her own ends:

Scenario 3 is identical to scenario 2, except that Sue is a cannibal who has a taste for human flesh and intends to ingest the arms of Kat, Liz, and Mel, unless Ben cuts off Jennifer's arms and serves them up to Sue's friend, Cannibal Tom.⁶⁰

We have in scenario 3 a case in which Jennifer cannot reasonably reject the minimizing-violation-permitting rule, since Kat, Liz, and Mel face the same fate—use as means only. Of course, Jennifer is faced with misuse at Ben and Tom's hands, while Kat, Liz, and Mel face misuse at Sue's hands. But this cannot make a difference with regard to the reasonableness of their rejecting the rules in question on the grounds that the harm constitutes a misuse. Misuse is misuse, regardless of who does the misusing.⁶¹

D. The Principle of Permissible Harm

Instead, the contractualist might try to employ the nonmoral factors identified by F. M. Kamm in her important invention, the Principle of Permissible Harm (PPH). This principle states that it is permissible to cause harm in order to produce a greater good just in case the harm

60. Maiming a human being because one enjoys the taste of his flesh and prefers to ingest it (rather than, say, readily available rice and beans) is an act that violates the Kantian dictum, if anything does. I take no position on the ethics of survival situations.

61. This is not to say that Ben has no reason to care whether it is he, or someone else, who does the misusing.

is an effect or aspect of the greater good itself that one produces. One may do (greater) good even though (lesser) harm will come of it. One may not, however, do harm in order that good will come of it (as a causal effect of the harmful act).⁶²

The Principle of Permissible Harm applies to our scenarios as follows. Ben's maiming of Jennifer has the causal effect that Kat, Liz, and Mel survive intact. But Jennifer's being maimed is not an effect or aspect of their survival intact. Therefore, PPH forbids Ben to maim Jennifer.

If PPH is valid then it seems to support nonminimizing restrictions. Less clear is how a contractualist can use the factors identified by PPH to support restrictions within his own framework. We know that Jennifer's being maimed has the causal effect that Kat, Liz, and Mel survive intact and that Jennifer's being maimed is not an effect or aspect of their survival intact. Let us grant the contractualist that these facts give Jennifer a special reason to reject the minimizing-violation-permitting rule. But we also know that the prospective maimings of Kat, Liz, and Mel are not effects or aspects of any greater good. In fact, when they are maimed no greater good is achieved. Their maimings merely have the effect that Sue achieves her hideous goal—no greater good, to be sure. It is doubtful that these patients should care a whit about the factors identified by PPH. They would seem to have at least as strong a reason to reject the nonminimizing-restriction rule as Jennifer has to reject the minimizing-violation-permitting rule. The factors identified by PPH do not give Jennifer a stronger reason to reject the minimizing-violation-permitting rule than Kat, Liz, or Mel has to reject the opposite rule. Therefore, Jennifer cannot reasonably reject the minimizing-violation-permitting rule. Once again, the contractualist's singular focus on the patient's perspective neutralizes potentially useful factors.

E. Justified Exceptions

A recent contribution by Scanlon suggests that he would analyze our scenarios in terms of justified exceptions to duties. Ben has a duty not to maim nonthreatening, nonconsenting, innocent individuals, such as Jennifer. Jennifer can reasonably reject any rule that permits Ben to maim her, unless circumstances justify an exception to that duty. The need to prevent Sue from causing greater harm to Kat, Liz, and Mel does not justify an exception to that duty, so Jennifer has reason to reject a minimizing-violation-permitting rule.⁶³

62. Kamm, *Morality, Mortality*, vol. 2, chap. 7.

63. I adapt Scanlon's treatment of the classic transplant case, in which the agent must decide whether to harvest organs from a healthy patient in order to save several patients with organ failure. Scanlon, "Intention and Permissibility," p. 313; Foot, "The Problem of Abortion and the Doctrine of Double Effect."

However, this account has not yet considered the perspectives of Kat, Liz, and Mel. The contractualist still owes us an explanation why the need to protect Kat, Liz, and Mel from Sue fails to justify an exception to Ben's duty. There are several possible lines a contractualist could take.

Scanlon himself argues as follows. He recognizes that Kat, Liz, and Mel would benefit from Jennifer's maiming, but he asserts that "to recognize this benefit as a ground for insisting on exceptions to principles protecting our [bodies] would be to take the view that our bodies are resources on which everyone has, at the outset at least, the same claim. It would thus be inconsistent with the idea that each person has a special claim to and authority over his or her own life and body." Scanlon takes this idea to be "one that no one could reasonably reject, and thus to be a notion that shapes other contractualist arguments."⁶⁴

However, if each person has a "special claim to and authority over his or her own body," then Kat, Liz, and Mel each have a special claim to and authority over his or her body, just as forceful as Jennifer's claim over hers. Scanlon has not explained why Jennifer's "special claim" is more important than three other such claims.

Scanlon might suggest that only Jennifer is in danger of being maimed by Ben, such that only her special claim against maiming is relevant for him in this case. Indeed, Ben has strong agent-relative reasons not to maim Jennifer. However, Jennifer's special claim cannot directly supply Ben with moral reasons, on Scanlon's view. It can only supply him with moral reasons indirectly, insofar as Jennifer's claim bears on the reasonability of rejecting a rule that permits minimizing violations. The question is not whether Ben has reasons against maiming Jennifer that no one can reasonably reject. It is whether anyone could reasonably reject a rule which permits Ben to maim Jennifer. Perhaps someone could reasonably reject a rule if it forbade him to act on his agent-relative reasons. But no one can reasonably reject a rule just because it permits one to treat one's agent-relative reasons as less than decisive. Such a rule does not limit one's ability to act on one's agent-relative reasons.

One might worry that the preceding argument begs the question by viewing from an agent-neutral standpoint special claims not to be treated only from an agent-neutral standpoint. This is not so. My argument does not simply view special claims from the agent-neutral standpoint. Rather, I argue that special claims do not have the requisite moral significance from within either the impersonal or the impartial standpoint, as Scanlon has constructed it. A theory that constructs its moral standpoint out of agent-relative standpoints in addition to that of the

64. Scanlon, "Intention and Permissibility," p. 313. I substitute my characters.

agent himself will share many features with the impersonal standpoint which proves so inhospitable to nonminimizing restrictions. Ultimately, both consequentialism and contractualism construct their respective moral standpoints out of agent-relative standpoints in addition to that of the agent. That is what makes them both patient-focused, and that is what prevents them from accommodating nonminimizing restrictions.⁶⁵

F. Inviolability

Scanlon might now advert to his thesis that “the general authorization or prohibition of a class of actions can have significance that goes beyond the consequences of the actions that are performed or not performed as a result.”⁶⁶ He could propose, following Kamm and Nagel, that general public acceptance of the nonminimizing-restriction rule confers *inviolable status* on Kat, Liz, and Mel, even as its acceptance by Ben leads Sue to violate these three.⁶⁷

There are several problems with the inviolability approach. First, it is not clear why public acceptance of the nonminimizing-restriction rule confers inviolable status (or any status whatsoever) on Kat, in particular. Our acceptance of the rule would seem, rather, to confer that status on Jennifer, protecting her as it does, while withholding that status from Kat, Liz, and Mel.

The defender of nonminimizing restrictions could respond by pointing out that the nonminimizing-restriction rule contains no proper names. It is entirely general. So we should understand acceptance of this rule to confer inviolable status, generally, on all human beings. Public acceptance of this rule confers inviolable status even on those, such as Kat, who are violated as a result.

It is true that the nonminimizing-restriction rule contains no proper names. We could interpret it as conferring inviolable status on all of humanity. But we could interpret the rule, just as sensibly, as conferring inviolable status on a certain subcategory of human beings—namely, those who are lucky enough to find themselves in Jennifer’s position rather than Kat’s. I see no reason to suppose that public adoption of a nonminimizing-restriction rule confers more inviolability on human be-

65. The thoughts contained in this paragraph, and the two that precede it, developed in response to comments by Paul Hurley, delivered at the Pacific Division of the American Philosophical Association in April 2003.

66. Scanlon, *What We Owe to Each Other*, p. 203.

67. Kamm, *Morality, Mortality*, vol. 2, pp. 271–89; Nagel, *Equality and Partiality*, p. 148.

ings, as a class, than would the adoption of certain minimizing-violation-permitting rules (a point I shall develop in the next section).⁶⁸

G. Expressing Exclusive Control

Perhaps the contractualist could strengthen his appeal to inviolability by incorporating more explicitly expressive considerations—those concerning the distinctive messages we publicly communicate when we accept one rule rather than another. Expressive considerations have a natural attraction for restrictionists. Public acceptance of any rule that permits maiming an innocent person does, indeed, express a message about the extent to which each of us enjoys exclusive control over his or her body. The restrictionist can insist that this message is simply intolerable. Of course, public rejection of the rule may have various unwelcome effects. But, so long as these are nonexpressive effects, the expressivist can insist that they make little, if any, normative difference. To this extent, an expressivist can support a nonminimizing-restriction rule, despite the fact that public acceptance of such a rule leads to the maiming of Kat, Liz, and Mel.

There are, in fact, foundationally expressive theories that hold out some promise for supporting nonminimizing restrictions. I shall argue, however, that expressivist considerations cannot help contractualists in this regard. Expressive considerations can support nonminimizing restrictions only when integrated into an agent-focused theory. The focus, I shall claim, has to remain on what the agent expresses by acting on a certain maxim, not what the society expresses by its adoption of the maxim. Elizabeth Anderson's expressive pluralism is one such theory. It includes a defense of strong deontic restrictions, and it is explicitly and foundationally agent focused.⁶⁹

68. Indeed, these observations suggest that nonminimizing restrictions may actually embody a certain bias. Most readers of this article believe that circumstances will never require them deliberately to attack an innocent person in order to secure their basic needs. By contrast, destitute individuals (e.g., in the Third World) lack necessities and often cannot obtain them except by attacking privileged persons (and then only if the opportunity presents itself). Suppose the destitute can protect their children from aggression only by stealing (in order to purchase defensive equipment or to relocate their children to safer regions). Anything weaker than nonminimizing restrictions would permit such robberies, under certain circumstances. Perhaps public acceptance of a nonminimizing-restriction rule confers inviolable status on privileged individuals, while denying that status to the destitute persons of the world. Compare Elizabeth Ashford, "The Demandingness of Scanlon's Contractualism," *Ethics* 113 (2003): 273–302. To this extent, such a rule would reflect what Scanlon once called "the bias of the lucky against the unlucky." T. M. Scanlon, "Rights, Goals and Fairness," in *Public and Private Morality*, ed. Stuart Hampshire (Cambridge: Cambridge University Press, 1978), p. 109.

69. See, e.g., Anderson, pp. 73–79. Anderson is a fellow impartialist, but her conception of impartiality is more radically pluralist than Scanlon's. Scanlon rejects the agent-centered aspect of Anderson's view (*What We Owe to Each Other*, pp. 384–85, n. 21).

By contrast, contractualism is an inherently patient-focused theory. An expressivist contractualism evaluates an action in terms of what a society would express by adopting a public rule which permits the action. Moreover, the question as to what the society expresses must be answered from the perspective of those who will be affected by the adoption of the rule. This is just to say that the patient's perspective remains decisive.

It is therefore important for contractualists who wish to defend restrictions not to be misled into supposing that they can simply appropriate the agent-focused virtues of foundationally expressivist theories such as Anderson's. Kumar may have made this mistake in his interesting contribution. Kumar asks us to consider "what is being *said* about persons by its being morally permissible for individuals to think of using another as a means (for a worthy end) without seeking her consent."⁷⁰ To this extent, Kumar incorporates into his contractualist framework an explicitly expressivist element. He asks us to consider "the way an individual will see herself as a result of the relations between individuals that result from a moral system in which it is permissible for a person to *commandeer* another's body in order to save her life."⁷¹ He contends that the minimizing-violation-permitting rule would undermine "each person's sense of herself as an independent being, whose body is *hers*, to the extent that she has sole authority in deciding how it will be used, abused, or developed."⁷²

To apply this thought to our hypotheticals, a minimizing-violation-permitting rule explicitly gives Ben the "authority" to maim Jennifer for the sake of Kat, Liz, and Mel—to "commandeer" her body. A minimizing-violation-permitting rule actually authorizes Ben to sever Jennifer's arm and deliver it to Tom. Consider, by contrast, how things stand under the restriction rule. Under the restriction rule, Sue maims Kat, Liz, and Mel, but that rule does not give anyone the authority to commandeer anyone else's body. When Sue maims Kat, she violates the restriction rule. For this reason, Kumar would argue, no one loses "her sense of herself as an independent being, whose body is *hers*."⁷³

The contractualist needs to argue that one cannot reasonably reject a rule that leads one to be maimed, in fact, but does not authorize one's maiming, in favor of a rule that authorizes one's maiming, in principle, but does not lead one to be maimed, in fact. The basic idea is that, for

70. Kumar, "Defending the Moral Moderate," p. 307, italics added.

71. *Ibid.*, p. 305. Quinn and Kamm make similar points. See Quinn, pp. 309–11; Kamm, *Morality, Mortality*, vol. 2, chaps. 9, 10.

72. Kumar, "Defending the Moral Moderate," p. 306, italics in original. Scanlon makes the same point, describing "the idea that each person has a special claim to and authority over his or her own life and body" as "one that no one could reasonably reject" ("Intention and Permissibility," p. 313).

73. Kumar, "Defending the Moral Moderate," p. 306, italics in original.

the purpose of applying the contractualist test, only some of the harms that result from the acceptance of a rule, R, count against R. The restrictionist can distinguish between (1) harms inflicted by an agent, acting pursuant to R and (2) harms that are not inflicted by anyone pursuant to R, even if they occur as a causal result of someone's accepting R. The restrictionist can then stipulate that harm counts against R only if someone inflicts it pursuant to R.

Suppose Ben acts pursuant to the minimizing-violation-permitting rule and maims Jennifer. According to the suggested criterion, Jennifer's maiming counts against the minimizing-violation-permitting rule because Ben maims her pursuant to the rule. By contrast, suppose Ben acts pursuant to the nonminimizing-restriction rule and refuses to maim Jennifer. Sue reacts by maiming Kat, Liz, and Mel. According to the criterion, these maimings do not count against the nonminimizing-restriction rule, even though they would not have occurred without Ben's adherence to that rule, because Sue herself does not maim pursuant to the nonminimizing-restriction rule. Sue accepts, instead, a minimizing-violation-permitting rule that authorizes her to maim.

Notice that, if we accept the nonminimizing-restriction rule, we are not rationally committed to regarding Kat's maiming as morally justified, though we know it will occur. Whereas, if we accept the minimizing-violation-permitting rule, then we commit ourselves to regarding Jennifer's maiming as justified, more like a consensual amputation or a maiming properly inflicted in self-defense.⁷⁴ To this extent, the defender of nonminimizing restrictions could argue that Jennifer has a stronger generic reason to reject the minimizing-violation-permitting rule than Kat has to reject the restriction rule, because the former rule commits us to regard Jennifer's maiming as justified, whereas the minimizing-violation-permitting rule does not commit us to regard Kat's maiming as justified. Jennifer can reasonably reject the minimizing-violation-permitting rule, while no one can reasonably reject the restriction rule.⁷⁵

The challenge for contractualists is to explain why Jennifer has a stronger agent-relative reason to reject a rule that authorizes her maiming than Kat has to reject a rule that, when obeyed, causes someone to maim her pursuant to an authorizing rule. If Ben maims Jennifer then Jennifer gets maimed pursuant to an authorizing rule. If Ben refuses to maim Jennifer then Kat gets maimed pursuant to an authorizing rule. This is not to deny that Kat might have an agent-relative reason to want Ben to adopt the nonminimizing-restriction rule. But so far it seems

74. Compare a related discussion of "independence" in Lippert-Rasmussen, "The Impermissibility of Minimizing Violations," pp. 345–46.

75. See related discussions in Kamm, "Harming Some to Save Others," pp. 251–55; and Rosenberg.

that Kat has a stronger agent-relative reason to want him to adopt the minimizing-violation-permitting rule.

Whether anyone can reasonably reject these rules depends on the point in time when the question is asked. Ex post, Jennifer can reasonably reject the minimizing-violation-permitting rule because she knows that it authorizes Ben to maim her, while the restriction rule does not authorize any maimings. Ex ante, however, matters are different. No one knows what role, if any, he will play in the scenario. I offer the following comparison. Consider, first, the extent to which one's sense of bodily independence would be compromised by the ex ante knowledge that one will be maimed if one finds oneself in the position of Kat, Liz, or Mel and someone in Ben's position obeys the moral rules accepted by one's society. Compare this degree of damage to one's sense of independence with the degree of damage that would result from the knowledge that one will be maimed if one finds oneself in Jennifer's position and someone in Ben's position proceeds to obey the moral rules accepted by one's society.

For Scanlon, the ex ante perspective seems to be the relevant one.⁷⁶ Suppose everyone has an equal chance, ex ante, of playing any given role. In that case, I think one's sense of bodily independence is at least as compromised by the nonminimizing-restriction rule as by the minimizing-violation-permitting rule. It seems to me that knowing that one might be maimed, simpliciter, compromises one's sense of independence to a very great degree. It is not clear that knowing that one might be maimed in Jennifer's position compromises one's sense of independence to any greater degree than does knowing that one might be maimed in Kat's, Liz's, or Mel's position. One's sense of independence is compromised, in either of these situations, because one knows that one is vulnerable to uninvited maiming. The fact that the minimizing-violation-permitting rule actually authorizes someone in Ben's position to maim someone in Jennifer's position does not seem to detract any further from Jennifer's sense of independence, or ours, ex ante.

Sue's adoption of a nonminimizing-restriction rule, in itself, seems to make the same statement as Ben's. We could even argue that, between Ben and Sue, Sue's adoption of a nonminimizing-restriction rule would

76. Scanlon writes that "in deciding which systems of principles are 'acceptable,' we cannot envision the reaction of every actual person. We can consider only representative cases" (*What We Owe to Each Other*, p. 171). The reaction of an "actual person" who has been affected by the public adoption of a certain principle is, by definition, an ex post perspective. So when Scanlon asserts that we *cannot* envision the reaction of every actual person he is asserting that we cannot take an ex post perspective. See also pp. 202–6, esp. p. 204 ("An assessment of the rejectability of a principle must take into account the consequences of its acceptance in general, not merely in a particular case that we may be concerned with").

confer more inviolability on human beings than would Ben's, since Sue's adoption leads to fewer violations. Perhaps, ironically, the nonminimizing-restriction rule even withdraws an important moral status from all human beings, including Jennifer, when Ben's acceptance of it causes Kat, Liz, and Mel to be maimed.⁷⁷ There may be reasons to favor one interpretation over the other, but I doubt that considering the issue within a contractualist framework reveals one interpretation to be uniquely plausible, or even more plausible than the alternative.

My point is that each of us has a strong, agent-relative reason to avoid being maimed, period. I do not see that one has any stronger reason to avoid being maimed by someone who acts on an authorizing moral maxim (i.e., someone who believes himself to be morally justified). Given that strange choice, one might even prefer to be maimed by someone who acts on an authorizing moral maxim. In any case, I see no reason why one should prefer the opposite.

To defend restrictions, the expressivist must argue that public acceptance of a minimizing-violation-permitting rule sends a more negative message than does acceptance of a nonminimizing-restriction rule. The contractualist, in turn, must explain this comparative judgment as follows. The minimizing-violation-permitting rule permits someone to maim me but reduces by a factor of (at least) three my ex ante chances of being maimed. By contrast, the nonminimizing-restriction rule forbids anyone to maim me but (at least) triples my ex ante chances of being maimed. The contractualist must insist that, given these two possible messages, I could not reasonably reject the nonminimizing-restriction rule in favor of the minimizing-violation-permitting rule.

1. Accommodating other factors.—I anticipate great difficulties for the contractualist who wishes to make this case. If she wants to invoke expressive considerations in her defense of nonminimizing restrictions, she must insist that the expressive factors associated with public rejection of such restrictions boast something like lexical priority over other factors relevant to reasonable rejection.

Consider the fact that commonsense morality permits defensive aggression. This includes, at the least, maiming an unjustified aggressor if this is necessary in order to defend oneself or another against his aggression. Even nonminimizing-restriction rules permit defensive aggression. Scanlon, too, appears to favor it.⁷⁸ Thus, it is revealing that the commitments which expressive contractualists must undertake in

77. Consider, e.g., the status enjoyed by beings who are so precious that it is morally permissible to sacrifice the one for the sake of saving the many. See Shelly Kagan, "Replies to My Critics," *Philosophy and Phenomenological Research* 51 (1991): 919–28, p. 920.

78. "One may use deadly force when this seems the only defense against a person who threatens one's life" (*What We Owe to Each Other*, p. 199).

order to defend nonminimizing restrictions render defensive aggression difficult (perhaps impossible) to justify within their framework. The expressive contractualist favors nonminimizing-restriction rules over minimizing-violation-permitting rules because the former express more forcefully the message that each individual enjoys exclusive control over her respective body. The nonminimizing-restriction rule states, "Do not maim innocent individuals." More fully, it states, "Do not maim any individual, X, unless either (1) X consents or (2) X threatens someone else with unjustified aggression, and maiming X is necessary in order to protect the other."

I shall not deny that this rule expresses, to some degree, the exclusive control message. But we can simplify the rule so as to express this message still more forcefully. We can eliminate the second qualifying clause, leaving us with the

Pacifist rule: Do not maim any individual unless she consents.

The pacifist rule forbids defensive aggression. Most of us reject it. But I fail to see how an expressive contractualist who supports nonminimizing restrictions could consistently fail to favor it, much less oppose it. When society adopts the pacifist rule it seems to express very forcefully the message that each individual enjoys exclusive control over her own body. Indeed, a society in which everyone actually adopts the pacifist rule seems much more appealing than our society.

Of course, in the real world, it is very likely that more unjustified violence will occur in societies in which all decent persons adopt the pacifist rule. Others will continue to initiate unjustified aggression, and a society adopting the pacifist rule will not defend its members against one another. Surprisingly, an exclusively expressive contractualism should have no problem with this state of affairs. Of course, when an aggressor threatens me with violence, he expresses a message inconsistent with exclusive control. But society at large expresses no such message when the aggressor acts. Society does not authorize his aggression, it simply foresees that such aggression will occur and declines to intervene.

The expressive contractualist must argue that I can reasonably reject the pacifist rule and the minimizing-violation-permitting rule in favor of the nonminimizing-restriction rule. How can she do this? She might suggest that the pacifist rule expresses a certain kind of approval for initiating unjustified aggression. Of course, the pacifist rule expressly forbids unjustified aggression, but accepting that rule removes the major obstacle from those who would perpetrate such aggression. Perhaps it sends a different message: "It is open season on defenseless individuals." Hence, one could argue that this message detracts from the overall message publicly expressed by the accepted rules, despite the fact that

the pacifist rule expresses the exclusive control message more strongly than does the nonminimizing-restriction rule. It renders the overall message more rejectable.

The foregoing argument relies on a crucial equivocation. The expressive contractualist argues that, by forbidding defensive aggression, public acceptance of the pacifist rule expresses a kind of approval of unjustified aggression. But the pacifist rule expressly forbids such aggression. Nor do decent people intend to encourage unjustified aggression when they accept this rule. At most, they foresee that their acceptance of the rule will result in more aggression, and they allow it to occur.

By this same reasoning, it would follow that public acceptance of a nonminimizing-restriction rule expresses a kind of approval of unjustified aggression perpetrated by wicked agents, such as Sue. This would be so even though the rule expressly forbids such aggression. The public may not intend for agents such as Sue to maim anyone when it adopts the rule, but it foresees that such aggression will occur as a result and allows it to take place. If someone can reasonably reject the pacifist rule then someone can also reasonably reject nonminimizing-restriction rules.

2. *Inconsistency and false dichotomies.*—Scanlon asserts that treating the benefits to Kat, Liz, and Mel as reasons to make an exception to Ben's duty not to maim is "inconsistent" with the idea that each person has a special claim to and authority over his or her own body. But where is the alleged inconsistency? There is nothing literally inconsistent about accepting the minimizing-violation-permitting rule along with the idea that each person has a special claim to and authority over his or her body. I can accept this rule while still believing, for example, that I have no right to maim you for my own amusement, or even to cut your hair without your permission. I can believe, in other words, that each person has a special claim to and authority over his or her body. What I cannot consistently believe, if I accept the minimizing-violation-permitting rule, is that each person has an absolute (or even nearly absolute) claim to and authority over her body (even if she is innocent, unconsenting, and poses no threats). Scanlon correctly observes that no one can reasonably reject the idea of the special claim. But perhaps someone can reasonably reject the idea of the absolute, or nearly absolute, claim. Indeed, to insist that no one can reasonably reject the absolute claim is to beg the question in favor of nonminimizing restrictions.⁷⁹

79. As Kamm observes, reviewing Scanlon's book, "Scanlon helps himself to many distinctions, at least analogous to intending versus foreseeing harm, and perhaps harming versus not aiding, whose moral relevance Standard Deontologists spend a lot of time trying to justify against consequentialist criticism. One would have liked to see how his contrac-

Kumar, similarly, states that a minimizing-violation-permitting rule makes it “morally permissible for individuals to think of using one another as a means (for a worthy end) without seeking her consent.”⁸⁰ This statement is true, but misleadingly incomplete. The minimizing-violation-permitting rule does make it permissible for individuals, in certain very special situations, to think of using someone as means without seeking her consent. This is different from making it permissible, generally, for individuals to think of using one another as means. Even under the minimizing-violation-permitting rule, I am not permitted to maim others just because I have a taste for human flesh and some good recipes! A practice of widespread homicidal cannibalism does not promote the good, in realistic worlds. Indeed, most practices that deontologists condemn as “using others as means only” do not ordinarily promote the good, so even act consequentialism forbids them, in most situations. A defender of an optimization-permitting axiom could characterize her statement about persons as a general statement, with some important exceptions. In general, she says, persons are to be revered, but exceptional cases arise in which that reverence must give way. The statement we make by accepting the minimizing-violation-permitting rule is not necessarily anything like the broad “insult to persons” that Kumar portrays it to be.

Restrictionists are led astray if they artificially limit the choice of rules to the simplest or most obvious. Consider a choice between a nonminimizing-restriction rule and a minimizing-violation-permitting rule that permits maiming for a range of utilitarian (or, worse, self-serving) purposes, such as:

Cannibals’ rule: Never maim innocent human beings, unless necessary to procure human flesh for your favorite recipes.

We can, to be sure, reasonably reject cannibals’ rule in favor of a nonminimizing-restriction rule. But there are other options, still more attractive. Consider the following minimizing-violation-permitting rule:

Recursive rule: Never maim innocents unless necessary to promote general acceptance of this rule.

The recursive rule is not a nonminimizing-restriction rule. But note the intuitive moral difference between the recursive rule and cannibals’ rule. I suggest that, if widely accepted, the recursive rule would express

tualism could help justify the moral relevance of ‘how’ factors rather than just taking it for granted. But he does not do this. . . . Much of Standard Deontology is not only not being justified, it is being assumed correct” (“Owing, Justifying, and Rejecting,” p. 340).

80. Kumar, “Defending the Moral Moderate,” p. 307.

the exclusive control message at least as effectively as would any alternative rule of which I am aware.⁸¹

Now consider this final hypothetical:

Scenario 4 is identical to scenario 3, except that Sue, along with most people, accepts cannibals' rule. Sue is ready to maim Kat, Liz, and Mel for cannibalistic purposes. However, along with most people, Sue desperately wants Ben to maim Jennifer. If Ben maims Jennifer, then Sue, and many others, will permanently abandon cannibals' rule and convert to the recursive rule, sparing Kat, Liz, and Mel. Ben understands all of this.

Scenario 4 equalizes Jennifer's situation with those of Kat, Liz, and Mel. By positively maiming, Ben can minimize the total number of positive maimings inflicted pursuant to a minimizing-violation-permitting rule (cannibals' rule, in this case). The only rule that would permit Ben to regard maiming Jennifer as morally forbidden, even under these circumstances, would be a nonminimizing-restriction rule. No optimization-permitting axiom can support such a restrictive rule. The closest an optimization-permitting axiom can come is to support a version of the recursive rule. Corollaries of this rule include not only "You may not positively maim if a practice of maiming under such circumstances would not minimize authorized positive maimings in the long run" but also "You may positively maim if a practice of maiming under such circumstances would minimize authorized positive maimings in the long run."⁸² The recursive rule does, under certain conditions, authorize positive maiming. It embodies only positive-harm-minimizing restrictions, not nonminimizing restrictions.

Can a contractualist instruct Ben to adopt a nonminimizing-restriction rule? Consider Jennifer's perspective. If Ben adopts the recursive rule and maims her, she thereafter enjoys a society in which most people adopt the recursive rule. Now consider Kat. If Ben adopts the nonminimizing-restriction rule then Sue maims Kat, and Kat continues to endure a society in which most people retain the cannibals' rule. Who has the greater objection? At best, Jennifer and Kat's objections seem evenly matched. So contractualism cannot support the nonminimizing-

81. Scanlon recognizes the "pressure toward making principles more fine-grained, to take account of more and more specific variations in needs and circumstances." But he emphasizes the "counterpressure arising from the fact that finer-grained principles will create more uncertainty and require those in other positions to gather more information in order to know what a principle gives to and requires of them" (*What We Owe to Each Other*, p. 205). He could argue that the recursive rule is too "fine-grained." Perhaps it creates too much uncertainty. If so, the defender of optimization-permitting axioms is just as entitled to appeal to this fact in her defense of nonminimizing restrictions.

82. See Amartya Sen, "Rights and Agency," *Philosophy of Public Affairs* 11 (1982): 3–39; Scheffler, *The Rejection of Consequentialism*; Nozick, *Anarchy, State and Utopia*.

restriction rule over the recursive rule.⁸³ In fact, I think Kat's complaint trumps Jennifer's, since Kat (and Jennifer) will remain vulnerable to untold additional maimings if Ben refuses to maim Jennifer. So I think Kat can reasonably reject the nonminimizing-restriction rule in favor of the recursive rule. If either of these conclusions is correct then contractualism, even when enhanced with expressivism, cannot support nonminimizing restrictions. Only from Ben's agential standpoint could Ben's personal acceptance of a minimizing-violation-permitting rule possibly appear worse than many other people's acceptance of that rule.

3. *The role of the partial compliance assumption.*—The preceding argument depends on the premise that a theory must deliver satisfactory results under conditions of partial compliance, rather than full compliance.⁸⁴ The contractualist could object that her theory is not addressed to partial compliance scenarios at all. She could insist that no one could reasonably reject the nonminimizing-restriction rule in favor of the recursive rule, if they were confident that everyone would adopt the same set of rules. However, I see no reason to accept this claim. Under conditions of full compliance, I am unsure why anyone would prefer the one society to the other, since no maimings would take place in either society.

The expressive contractualist could argue that, under conditions of full compliance, it would be unreasonable to reject the rule that expresses the more exalted and peace-loving moral ideal (i.e., the nonminimizing-restriction rule). But this same reasoning backfires. If no one could reasonably reject the nonminimizing-restriction rule in favor of the recursive rule, under the full compliance assumption, then it seems that, likewise, no one could reasonably reject the pacifist rule in favor of the nonminimizing-restriction rule, under those conditions. If everyone accepts either the pacifist rule or the nonminimizing-restriction rule then no one maims, as a matter of fact. Yet public acceptance of the pacifist rule seems to express the exclusive control message even more forcefully than does the nonminimizing-restriction rule, which permits self-defense. How could it be reasonable, *ceteris paribus*, to

83. The contractualist cannot claim to support nonminimizing restrictions just by showing that the recursive rule is reasonably rejectable. She must also show that the nonminimizing-restriction rule is not reasonably rejectable. Otherwise, she will merely have shown that refusing to commit minimizing violations is morally permissible. The supporters of optimization-permitting axioms need not deny this.

84. Apparently other interpreters have had trouble discerning whether Scanlon intends contractualism as a full- or partial-compliance theory. Thomas Pogge writes of Scanlon's book, "My study of the text had led me to interpret Scanlon as favoring the full compliance stipulation. . . . But Scanlon has indicated, in correspondence, that he meant to stipulate partial compliance" ("What We Can Reasonably Reject," *Noûs* 11, suppl. [2001]: 118–47, p. 133).

reject a rule which allows us to bask in the glow of that exceptionally exalted message?

In the real world, of course, compliance is never more than partial. Innocents find themselves threatened with unjustified aggression. Non-minimizing-restriction rules permit us to use violence to protect one another from aggression. Therefore, a contractualist would insist, such rules enjoy an advantage over pacifist rules. This advantage “compensates” for the fact that nonminimizing-restriction rules express the exclusive control message less forcefully than do pacifist rules.

Likewise, the advantage of the recursive rule over nonminimizing-restriction rules only emerges in certain partial compliance situations. A nonminimizing-restriction rule forbids maiming, even when others threaten to violate the nonminimizing-restriction rule, whereas the recursive rule permits maiming if, but only if, others threaten to violate the recursive rule. If everyone adopts the nonminimizing-restriction rule, then no one maims. The same holds if everyone adopts the recursive rule. If everyone adopts the recursive rule, of course, then everyone reserves the right to maim. All things equal, we might prefer to live in a world in which no one reserves that right. But all things are not equal. Remember, all things equal, we would also prefer to live in a world in which no one reserves the right to maim aggressors, even in defense of others. In the real world, in which some people initiate unjustified aggression, however, we favor a society in which other agents (e.g., police officers) reserve the right to respond with defensive aggression. Similarly, all things equal, I would prefer to live in a world in which no one reserves the right to maim innocents. But, given the real world in which someone might maim me, I might reasonably prefer to live in a world in which others reserve the right to maim the innocent, when necessary to protect me from maiming at the hands of the wicked. I might rationally prefer this, even understanding that I might become the victim of this rule, rather than its beneficiary.

IV. CONCLUSION

There seem to be no factors that make it reasonable, from any standpoint, to reject a harm-minimizing-restriction rule or positive-harm-minimizing-restriction rule in favor of a nonminimizing-restriction rule. Contractualism gives us no reason to favor nonminimizing restrictions. In fact, it gives us some reason to favor harm-minimizing or positive-harm-minimizing over nonminimizing restrictions. I may have overlooked factors to which one could appeal in a contractualist defense of nonminimizing restrictions. But it falls now to contractualist-restrictionists to identify these factors.

Some restrictionists may have assumed that the restriction skeptics’ challenge was directed at all restrictions, including harm-minimizing

and positive-harm-minimizing restrictions. Much restrictionist ink has been spilled defending these weak restrictions against this fictional challenge. In the process, restrictionists have not noticed that contractualism offers nothing unique to the defense of weak restrictions. Some of the justificatory work in contractualist defenses of restrictions is being done by its structure as a two-level or indirect theory. The remainder of the work is done by deontic value theory. Both features are equally available to proponents of optimization-permitting axioms. What is worse, contractualists never get around to defending the nonminimizing restrictions that serve as the restriction skeptics' actual targets.⁸⁵

Scanlon once speculated that act utilitarianism would not emerge as a theorem of contractualism.⁸⁶ He was probably correct. But my conclusions suggest that contractualism could support, or be extensionally equivalent to, a different kind of optimization-permitting norm, such as indirect consequentialism or an indirect version of Scheffler's hybrid theory. Contractualism could support, as a corollary, an indirectly optimization-permitting norm, perhaps combined with a deontic value theory.⁸⁷ In turn, this corollary could support harm-minimizing or positive-harm-minimizing restrictions. But contractualism boasts no comparative advantage, as a foundation for restrictions, over optimization-permitting theories such as consequentialism.

85. My conclusions seem consistent with Richard J. Arneson's in "The End of Welfare As We Know It? Scanlon versus Welfarist Consequentialism," *Social Theory and Practice* 28 (2002): 315–36. I encountered Arneson's paper as mine was going to press. Derek Parfit makes some overlapping points about contractualism in an unpublished Tanner Lecture on Human Values which he delivered at the University of California, Berkeley, in December 2002. Some of these points appear in Parfit's contribution to Philip Stratton-Lake, ed., *On What We Owe to Each Other: Scanlon's Contractualism* (Oxford: Blackwell, 2003).

86. Scanlon, "Contractualism and Utilitarianism," p. 119. See also Nagel's discussion of R. M. Hare in *Equality and Partiality*, p. 44.

87. This paragraph interprets contractualism as a metaethical theory. If we understand contractualism as normative, then it could not support the mentioned optimization-permitting norms but could be extensionally equivalent to them (forbidding and permitting the same actions).