

The Problem of Rational Compliance with Rules

BEN EGGLESTON

Department of Philosophy, University of Kansas, 1445 Jayhawk Blvd., Lawrence, KS 66045, USA; e-mail: eggleston@ku.edu

1. Introduction

Let us suppose that the citizens of some country recognize the need for certain public goods, such as education, police, and roads. Different rules promoting the provision of such goods might have different levels of effectiveness, and they might vary in their costs of implementation or enforcement, too. For example, a rule leaving the provision of such goods up to private charity would probably have low costs but rate low on any reasonable measure of effectiveness, too. Let us suppose, then, that when all the costs and benefits are taken into account, the rule with the best consequences, morally speaking, requires the provision of public goods to be centrally managed, through taxation and appropriation. This optimal rule will presumably be fairly complex, including clauses concerned with the behavior of legislators, executive agencies, and ordinary citizens. Most of this complexity can be set aside, but it is important here to suppose that one aspect of this optimal rule requires every citizen to pay all the taxes that he or she legally owes.

Let us also suppose that a system in accord with such a rule gets started, with the rule being regarded as morally binding solely because of the consequentialist reasoning that can be offered in justification of it. With the system up and running, each year tax payments flow in and appropriations flow out, providing public goods in certain quantities. Inevitably, there is always some tax evasion, since not everyone is adequately motivated to comply with the optimal rule and the costs of increasing enforcement enough to catch every last evader would outweigh the benefits of securing those missing dollars. But the system is functioning, overall.

Finally, let us suppose that a particular citizen, aware of unmet needs in her local community, can see that if she were to follow the rule requiring her to pay all of her legally assigned taxes, the consequences would be worse, overall, than if she were to direct some or all of those

funds to addressing nearby unmet needs. She also reflects on the fact that the rationale for the rule is strictly consequentialist: it was adopted, and is generally observed, for the sake of promoting good consequences. If she were in such a situation knowing that compliance with the rule would have worse consequences than deviation from it would and knowing that the justification for the rule is strictly consequentialist, how could it be rational for her to opt for compliance anyway? The citizen's problem can be summed up more concisely, and more generally, as follows: How can it be rational for an agent to follow a rule with a purely consequentialist justification in a case in which she knows that she can do more good by breaking it?

This is the problem of rational compliance with rules. It is a perennial objection to most forms of rule consequentialism, and it remains a vigorously debated and unresolved problem in contemporary moral philosophy. As occasionally happens with long-unresolved philosophical problems, the specific contours of this problem are sometimes lost sight of, resulting in the failure of certain proposed solutions to address the precise problem at hand. We will consider two ways in which this can happen with the problem of rational compliance with rules, using Alan Goldman's recent article "The Rationality of Complying with Rules: Paradox Resolved" as a case study.¹ The first misconstrual of the problem is the more obvious of the two, and can be described and corrected briefly. The second misconstrual is more subtle, but also more interesting.

2. The Justification Gap

One way of losing sight of the contours of the problem of rational compliance with rules is to grant that the agent can do more good by breaking a rule than by following it, but to ignore the stipulation that the agent knows this, and to assume instead that the agent merely seems to know it, or merely believes it. On any standard approach to epistemology, this change in the problem removes the element of justification that implicitly attaches to the agent's belief when it is described unqualifiedly as a case of knowledge, as in the original problem.

In several passages, Goldman introduces just such a justification gap into the original problem. First, he writes that once the rule is justified consequentially, "[t]he problem remains to explain why it is rational for each to comply when she seems to know that she can do better by deviating, given that the effect of her deviation on the collective outcome is negligible."² Then, in outlining his case for compliance, he explicitly rejects the knowledge assumption of the original problem when he writes

that “individuals will not have a way of knowing how many other individuals are reasoning in the same way” as they are when contemplating nonpayment.³ Later in his article, he explicitly refers to belief, not knowledge, in writing “[o]ur question is why individuals ought to comply with rules that they regard as justified when they believe that they can do better by defecting in the individual case.”⁴ Again, epistemic concerns are the basis for Goldman’s claim that individuals “cannot coordinate to the optimal distribution of compliers and deviators.”⁵ A final telling remark to note here is Goldman’s claim that “there is no way to check that defectors remain below the threshold at which collective harm begins to set in.”⁶ Additional remarks could be cited to provide further confirmation of a justification gap in Goldman’s construal of the problem, but the foregoing remarks provide the most direct evidence.

Admittedly, Goldman also uses language that mentions or entails knowledge, and thus reflects the justification element of the original problem. For example, he writes: “One cannot rationally choose less than the best known option when the whole point is to optimize.”⁷ Additionally, he refers to the individual contemplating nonpayment as “realiz[ing]” that the consequences of nonpayment would be better than the consequences of compliance.⁸ Finally, he uses the phrase “known better alternative.”⁹ These are the only justification-acknowledging remarks that are really clear in their import. Moreover, they do not cancel the effect of the contrary remarks mentioned previously and the contexts in which they appear. That effect, clearly, is to modify the original problem by introducing a justification gap into it. Indeed the same effect is present, to a certain extent, in Goldman’s discussion of rules in his book *Practical Rules: When We Need Them and When We Don’t*.¹⁰

When the original problem is modified in this way, the result is an altered problem of compliance that can be solved straightforwardly. The fact that the rule with the best consequences requires compliance creates a presumption in favor of the claim that the consequences of compliance are usually better than the consequences of deviation. That claim, in turn, creates a presumption in favor of the claim that, in the agent’s particular case, the consequences of compliance are better than the consequences of deviation. This last claim may turn out to be false; indeed it is stipulated to be false. But if the agent is not justified in believing that, then she can rationally go along with the presumption in favor of the claim that, in the agent’s particular case, the consequences of compliance are better than the consequences of deviation. Thus, it can be rational for her to comply.

This account is a non-starter as a solution to the original problem, since the original problem contains the stipulation that the agent knows that the consequences of compliance are worse than the consequences of

deviation. As a result, introducing the justification gap into the original problem results in a substantially different problem that invites a reply that does not contribute to solving the original problem. To be sure, the original problem is not all that matters, philosophically. Introducing the justification gap could be a fruitful alteration of the original problem if, while not contributing to solving the original problem, it were to result in a more interesting problem. Unfortunately, the justification gap does not do this, either. The altered problem is uninteresting, because its solution is so easy. Plainly, it can be rational for an agent to follow a rule when she has a presumptive reason to do so and no justification for believing anything to the contrary.

Admittedly, the altered problem is not completely lacking in philosophical interest, in part because the easy solution given above could be challenged on various grounds. For example, a proponent of an objectively oriented conception of rationality concerned with actual consequences rather than foreseeable consequences could argue that it is irrational to choose an act whose consequences are sub-optimal, even if one has no way of knowing that they are sub-optimal. Such a view could turn out to be unexpectedly plausible. If it did, then the altered problem with the justification gap would turn out to be unexpectedly interesting. But that possibility is highly speculative, and the original problem does not require any such machinations in order to be philosophically interesting. Thus, modifying the original problem by introducing the justification gap neither contributes to solving the original problem nor opens other vistas of comparable philosophical significance.

3. Unfairness as an Underappreciated Bad Consequence of Deviation

Goldman's introduction of the justification gap into the original problem is only part of his response to it, and is secondary to a more interesting and subtle account that does not depend on the justification gap. This account involves working with a variant formulation of the problem of rational compliance with rules that includes a proviso distinguishing it from the original problem, and involves offering an argument for the rationality of compliance that hinges on this proviso. This particular strategy of Goldman's is an example of a family of responses that all misconstrue the original problem in essentially the same way, but it will be most instructive to focus on the details of Goldman's account before drawing that more general lesson.

Goldman presents his account in the context of a tax scenario of the kind discussed above. In order to understand his account correctly, it is

useful to follow him in avoiding some possible misunderstandings of this central example. First, it must not be thought that self-interest is the source of the agent's temptation to deviate. The agent is not trying to avoid parting with the prescribed sum of money; instead, she believes in the good that the rule was adopted to promote and just sees a better way to promote that good than by following the rule. She has "moral reasons for deviating from morally justified rules."¹¹ Indeed, throughout, the kind of rationality under discussion is not self-interested rationality: it is just rationality about moral considerations conceived of essentially consequentially. Second, "the rule itself does not constitute a reason for following it"; instead, "only ordinary consequential reasoning . . . counts."¹² Finally, it must not be thought that cases such as the one under consideration arise only because the rule in question is less justified, on consequentialist grounds, than a rule that would allow the funds to be diverted in such cases. A rule allowing diversion, while attractive in principle, would be impractical in many cases, and this undermines any consequentialist justification that might be offered for it.¹³ Thus, the rule most justified in terms of its consequences might well require some individuals to act, on occasion, in ways that do not result in the best consequences.

With these aspects of the problem of rational compliance with rules in view, let us consider the details of Goldman's account, which hinges on the addition of a proviso to the original problem. Here, again, is the original problem of rational compliance with rules: How can it be rational for an agent to follow a rule with a purely consequentialist justification in a case in which she knows that she can do more good by breaking it? Here is the problem that results from the addition of Goldman's proviso: How can it be rational for an agent to follow a rule with a purely consequentialist justification in a case in which she knows that she can do more good by breaking it, at least when there are no adverse consequences due to the existence of the rule itself? What Goldman means by "adverse consequences due to the existence of the rule itself" can be made more concrete with an example. Let us suppose that one of the rules of a park requires Amy to keep her dog on a leash. If she lets her dog run loose there, and people see her doing that, then it might be the case that an adverse consequence of her action is that other people take all of the park rules less seriously. Were it not for the existence of the rule, this adverse consequence would not occur. For Amy to let her dog run loose might be somewhat harmful, but it would not have the adverse consequence of weakening respect for the rules of the park. This adverse consequence is the sort of thing Goldman means by an adverse consequence due to the existence of the rule itself.

The difference between the two problems is significant. The original problem concerns cases in which the consequences of breaking the rule are better, all things considered, than the consequences of following it. Such cases can usefully be called desirable-deviation cases, though it must be noted that this is just meant to signify that the consequences of deviation are better than the consequences of compliance, not to presume that compliance is irrational in such cases. The original problem concerns how, in such cases, it can be rational to follow the rule anyway. In contrast, the altered problem that results from the addition of Goldman's proviso concerns cases in which the consequences of deviation are better than the consequences of compliance, not when all things are considered, but when consequences due to the existence of the rule are not counted against breaking it. The altered problem concerns not only desirable-deviation cases, but also another set of cases. Desirable-deviation cases form a proper subset of the cases that the altered problem concerns. It is plain that it concerns desirable-deviation cases. If breaking the rule has better consequences, all things considered, than following it does, which is definitive of desirable-deviation cases, then *a fortiori* breaking the rule has better consequences when we disregard certain consequences that would count against breaking the rule: specifically, adverse consequences due to the existence of the rule. But what makes the difference between the two problems significant is that the altered problem also concerns another set of cases.

The other cases are not desirable-deviation cases, but would be desirable-deviation cases if one were to omit counting, against deviation, adverse consequences due to the existence of the rule. In such cases, the only thing stopping deviation from having better consequences is that the existence of the rule gives rise to consequences that tip the scales in favor of compliance. Returning to Amy and her dog, let us suppose that, due to the existence of the rule, the consequences of letting her dog run loose would be worse than the consequences of keeping her dog on a leash, say, because letting her dog run loose would cause people to take the rules of the park less seriously. Thus, this is not a desirable-deviation case. But let us suppose also that, if one were to omit counting the adverse consequences due to the existence of the rule, then the consequences of letting her dog run loose would be better than the consequences of keeping her dog on a leash. So described, this would be a desirable-deviation case, were it not for those adverse consequences due to the existence of the rule.

The altered problem concerns not only desirable-deviation cases, but also these other cases. The reason for this is Goldman's proviso: it is what distinguishes the altered problem from the original problem, which has to do with the rationality of compliance in desirable-deviation cases only.

Due to Goldman's proviso, the altered problem invites us to consider the rationality of compliance in an additional set of cases in which the existence of the rule makes deviation have adverse consequences that tip the scales in favor of compliance.

To solve the altered problem, Goldman offers the following argument. First, "The rationality of adopting the rule indicates that the group itself has an overriding obligation to achieve the collective outcome at which the rule aims."¹⁴ Second, "The obligation of the group entails an obligation on the part of its individual members to do their part to contribute to the fulfillment of the obligation."¹⁵ Finally, it is unfair for an agent not to meet her obligation to do her part when others are meeting their obligations to do theirs.¹⁶ Therefore, the existence of the rule generates a new consideration based on the concept of fairness that can make compliance have better consequences, even if deviation would have had better consequences in the absence of adverse consequences due to the existence of the rule.

It might seem that by appealing to the concept of fairness, Goldman is illicitly invoking a deontological notion, or a side constraint, and is thereby failing to provide a genuinely consequentialist rationale for compliance. But Goldman neatly forestalls this charge by casting the unfairness of deviation as a bad consequence that goes into the scales with all the other consequences of the act, rather than as a non-consequentialist consideration meant to trump consequentialist considerations.¹⁷ By construing unfairness as just one consequence among many, Goldman is compelled to allow that this consideration is not always strong enough to outweigh the benefits of deviation.¹⁸ Oddly, Goldman himself faults one account for providing reasons in favor of compliance that could "always be overridden by moral reasons on the side of deviation in the individual case."¹⁹ In any event, the fairness consideration may be strong enough to outweigh the otherwise-decisive benefits of deviation in many cases; the unfairness of breaking a rule that others are following may be a pretty bad consequence. Thus, the appeal to fairness offers a solution to the altered problem. The appeal to fairness shows that even when consequences excluding those based on the existence of the rule favor deviation, consequences including those based on the existence of the rule can favor compliance, thus providing a genuinely consequentialist rationale for compliance.

There are a number of concerns one might raise about the nuts and bolts of Goldman's account. First, one might ask whether a group is the sort of entity that can have an obligation, strictly speaking. Second, even granting the notion of group obligation, one might think that the problem of rational compliance with rules can arise in cases that are not morally weighty enough to be cases of group obligation, as opposed to pursuits that the group may adopt at its discretion. For example, it would seem to

be possible for cases to arise in which a public good is not important enough for the group to have an obligation to provide it, but is nonetheless a good that the group can rationally adopt a rule to pursue, and is a good that some agent might be better able to promote by breaking the rule than by following it. Since Goldman derives the agent's obligation to follow the rule from the group's obligation to provide the public good, his account would appear to offer no rationale for compliance in such cases. It might be thought a reply to this objection that Goldman acknowledges that not all cases of group obligation are cases of moral obligation, noting that they may be cases of rational obligation instead.²⁰ But even rational obligations are still obligations, not discretionary pursuits. Finally, one might doubt that an agent who breaks the rule to promote its objective more fully is, in fact, acting unfairly. After all, as Goldman allows, the agent is not "refusing to sacrifice self-interest for the common good."²¹ On the contrary, the agent is making a sacrifice of the required amount, and is doing more for the common good by deviating than by complying. With respect to the justification gap, if the agent were not justified in believing herself to be in such a situation, then she might reasonably be accused of unfairness or, at least, bad judgment; but if she is justified in believing herself to be in such a situation, according to however rigorous a standard of justification for belief Goldman would like to impose, then it would seem that she should be applauded for her alertness, initiative, and effectiveness in promoting the general good. While Goldman has more to say about the unfairness of deviating in such cases, this matter and the previous two may warrant further scrutiny.

4. The Strategy of Demanding a Recount

A further concern is more fundamental, since it points to an aspect of Goldman's account that puts it in a family of responses to the original problem that are mistaken in essentially the same way. The defining feature of this family of responses is that of stipulating, presuming, or fearing the exclusion of some consequences from the initial tally of consequences for and against compliance that may militate in favor of deviation and then pointing out that compliance may come out on top once the excluded consequences are included in the final tally. When the tally of consequences seems to favor deviation, the responses demand a recount, each urging that special attention be paid to its favored class of possibly neglected bad consequences of deviation or good consequences of compliance. Unfortunately, such responses miss the point of the original problem of rational compliance with rules.

Let us consider again a case giving rise to that problem, a desirable-deviation case, where one is in a situation in which one has to decide whether to follow a rule that has a purely consequentialist justification, and one has taken into account all of the consequences of following the rule and all of the consequences of breaking the rule, and one has found the consequences of breaking the rule to be better than the consequences of following it. If one then asks how it can be rational to follow the rule rather than break it, what does a recount-demanding response imply? Unfortunately, all it implies is that deviating from the rule might have a bad consequence that one failed to take into account or that complying with it might have a good consequence that one failed to take into account. But the question at hand presupposes that one has already taken all the consequences into account; being told that one might not have tallied them up correctly is, strictly speaking, beside the point. Instead of being answers to the question at hand, recount-demanding responses are just queries to the main premise that the consequences of deviation are better than the consequences of compliance.

Recount-demanding responses would have us rationalize compliance by showing that, despite appearances, the consequences of compliance really are better than the consequences of deviation, typically by highlighting a possibly underappreciated bad consequence of deviation. Goldman's account, with its argument about unfairness, is of this kind. Other responses belonging to this recount-demanding family could be generated by appealing to the bad example one might set by breaking the rule, or by appealing to additional harms that might result from breaking a rule that others are following. For example, although Goldman's account is mainly about unfairness, he also suggests that one's "[s]elf-respect and self-identity" could be compromised if one were to break a rule that others are following.²² Many more responses of this kind could be generated: they are as plentiful as the bad consequences that can be said to result from the breaking of a rule.

To give a sharper characterization of the failure of such responses to address the original problem, it is useful to contrast them with responses given in attempts to show that even if the consequences of compliance are worse than the consequences of deviation, compliance can still be rational. Adapting a term from the title of Goldman's paper, the responses can rightly be called paradox-addressing responses. Because they would have us rationalize compliance even when it has worse consequences, paradox-addressing responses are genuinely pertinent to the original problem in which it is assumed that the bad consequences of deviation have been properly appreciated. A recount-demanding response, in which there

is a concern with possibly underappreciated bad consequences of deviation, is no substitute for a paradox-addressing response.

This point becomes especially vivid when teaching students about utilitarianism and the objection that utilitarianism can require, for example, punishing an innocent person for some crime if better consequences would result from doing so than from continuing to search for the real criminal. In teaching students about this objection, the point is to get the students to think about how much weight should be given to consequentialist reasoning when consequentialists prescribe a course of action that many people would regard as unjust. An inevitable distraction to be overcome, before the class can be brought to really focus on that issue, comes from students who hurry to rescue utilitarianism by insisting that the consequences of punishing the innocent person would not actually be better for all of the usual reasons about public outrage and setting a bad example. In reply to such students, one gives the refrain that all ethics teachers know well: "Even if this is true in many cases, it does not address the deeper issue concerning the conflict, in principle, between consequentialist reasoning and deontological values such as justice." In the context of the original problem of whether compliance can be rational in desirable-deviation cases, recount-demanding responses, which include arguments about the bad consequences of deviation, are like the interventions of the students. They are worth thinking about to make sure the consequences have all been tallied correctly, but they are not truly applicable to the problem at hand.

Recount-demanding responses dodge the problem of rational compliance with rules rather than actually addressing it. Matters are further obscured when a recount-demanding response is put forward as if it could be held to accomplish what a paradox-addressing response could, if it were successful. In this regard, Goldman's account can again serve as a case study. Admittedly, in defense of Goldman against this charge of overstating the significance his account could have, it might be pointed out that he conspicuously modifies the original problem with a proviso that decisively influences his entire account. It might also be argued that when he claims to have shown "that we only seem to have overriding reasons" in favor of deviation in certain cases, he is acknowledging that his response is in the recount-demanding family.²³ But that remark also admits of a broader reading, one consistent with his seeing his account as a paradox-addressing response. Most tellingly of all, Goldman explicitly discusses three other responses to the original problem, claims that they are unsuccessful, and claims that his response succeeds where they fail. The crucial fact about the three accounts Goldman discusses is that they are paradox-addressing responses, not recount-demanding responses.

Goldman considers his account to be better than the accounts of Edward McClennen, David Gauthier, and Christopher McMahon. Goldman cites a paper in which McClennen asserts the central claim of his theory of resolute choice: that even when the consequences of deviation are better than the consequences of compliance, compliance is rational if the outcome of compliance will be better than the outcome that would have resulted if the rule had not been adopted at all.²⁴ In the cases McClennen discusses, deviation has better consequences all things considered. He writes that “the rule counsels one to choose in a manner that will not always ensure that one chooses in accordance with the balance of reasons that arise within the context of a particular act of choice.”²⁵ He also writes that “[b]ecause one still foregoes the additional gain that could be realized by deviating from the rule, the balance of reasons could still be said to support such a deviation.”²⁶ Furthermore, both of these passages occur in contexts in which McClennen explicitly rejects a position he calls compatibilism, the position that “a particular instance of rule-guided choice is rational only if that choice is supported by the balance of reasons that apply to that case.”²⁷

McClennen’s account is complicated by the fact that he sees committing oneself to a rule as causing a change in one’s preferences that makes compliance more preferred than it would be in the absence of the commitment.²⁸ This might appear to make McClennen’s account an example, as Goldman’s account is, of the strategy of showing the consequences of compliance to be better than the consequences of deviation, after all. But for McClennen, such preference changes do not make compliance supported by the balance of reasons that obtain in the particular context; they just enable compliance to be rational even when deviation remains supported by the balance of reasons that obtain in the particular context, as indicated in the passages above. In his account, McClennen takes the consequences of deviation to be better than the consequences of compliance all things considered and still tries to rationalize compliance. McClennen’s account is a paradox-addressing response, not a recount-demanding response. Indeed this orientation has been characteristic of McClennen’s body of work.²⁹

With respect to Gauthier, Goldman cites the sixth chapter of *Morals by Agreement*.³⁰ There Gauthier presents his theory of constrained maximization and maintains that compliance can be rational even when the consequences are worse than the consequences of deviation, as long as compliance is dictated by a disposition, such as a cooperative disposition, that is rational for the agent to have. As with McClennen, at issue are cases in which deviation has better consequences than compliance, even when all consequences, including any that depend on the existence of the

rule, have been taken into account.³¹ Like McClennen, Gauthier has consistently oriented his work on practical rationality to respond to the problem of rational compliance with rules with evolving versions of a paradox-addressing response, not a recount-demanding response.³²

Finally, Goldman's discussion of McMahan focuses on his recent book *Collective Rationality and Collective Reasoning*.³³ There McMahan defends a "Principle of Collective Rationality."³⁴ For our purposes, this principle is relevantly similar to what is found in the accounts of McClennen and Gauthier, as it includes conditions under which compliance is rational even when deviation would have better consequences, all things considered. The three authors in whose company Goldman places himself, and to whose accounts he regards his own as a superior rival, all offer paradox-addressing responses to the original problem. By framing his account in this way, Goldman suggests that it belongs in that same category.

To his credit, Goldman sharply distinguishes his strategy from the strategies of McClennen, Gauthier, and McMahan, by repudiating departures from the orthodox straightforwardly maximizing conception of rational choice and endorsing a strategy of arguing that certain apparent desirable-deviation cases are not actually such cases because of oversights in the tallying of consequences.³⁵ Goldman is arguably right to regard the accounts of McClennen, Gauthier, and McMahan, as unsuccessful. But there is another, far more important, difference between Goldman's account and theirs, which is that his account, if it were successful, would not result in what their accounts would result in, if they were successful. The general lesson to be drawn is important: recount-demanding responses must be distinguished from paradox-addressing responses, especially since advocates of recount-demanding responses may see and depict them as superior alternatives to paradox-addressing responses.

5. The Role of Recount-Demanding Responses

While recount-demanding responses must be kept distinct from paradox-addressing responses, since it is only with paradox-addressing responses that we may solve the problem of rational compliance with rules, recount-demanding responses have a role to play here. A recount-demanding response, when sound, is of value in helping to expose the true cost of deviation in cases in which rules exist. Moreover, in practice, a recount-demanding response, though it cannot be used to do what a paradox-addressing response is designed to do, might obviate the need for a paradox-addressing response to be given. Solving the problem of rational

compliance with rules or, as may be said, resolving the paradox, might not actually be necessary. If a person is trying to persuade someone else to follow a rule in an apparent desirable-deviation case, and the person has no solution to the problem of rational compliance with rules, then the person might be able to secure compliance anyway by showing that deviation has underappreciated costs that prevent the case at hand from really being a desirable-deviation case. As we saw earlier, this strategy might be pursued in any number of ways, whether by drawing on Goldman's argument about the unfairness of deviation or by appealing to other bad consequences of deviation, such as setting a bad example for others. Thus, sound recount-demanding responses should be credited with adding to the arsenal of considerations that we can adduce in order to establish that various apparent desirable-deviation cases are not genuine desirable-deviation cases. Such considerations show that desirable-deviation cases are rarer than they might sometimes be thought to be, and thereby lessen the urgency of showing that compliance in such cases can be rational. But lessening the urgency of solving the problem of rational compliance with rules is very different from actually solving it. Recount-demanding responses might help us to see that, in many cases, bringing about the best consequences does not require us to break a rule. But in cases in which bringing about the best consequences does require us to break a rule, such responses will not help us to know whether we should.³⁶

Notes

1. See Alan H. Goldman, "The Rationality of Complying with Rules: Paradox Resolved," *Ethics* 116 (2006).
2. *Ibid.*, p. 455.
3. *Ibid.*, p. 456.
4. *Ibid.*, p. 463.
5. *Ibid.*, p. 464.
6. *Ibid.*, p. 465.
7. *Ibid.*, p. 457.
8. *Ibid.*, p. 458.
9. *Ibid.*, p. 460.
10. Alan H. Goldman, *Practical Rules: When We Need Them and When We Don't* (Cambridge, England: Cambridge University Press, 2002). Also see Ben Eggleston, "Alan H. Goldman, *Practical Rules: When We Need Them and When We Don't*," *Utilitas* 16 (2004), esp. p. 115.
11. Goldman, "The Rationality of Complying with Rules: Paradox Resolved," *op. cit.*, p. 454.

12. Ibid., p. 455.
13. Ibid., pp. 455–456.
14. Ibid., p. 462.
15. Ibid., p. 463.
16. Ibid., p. 464.
17. Ibid., pp. 464–465.
18. Ibid., p. 468.
19. Ibid., p. 458.
20. Ibid., pp. 462–463.
21. Ibid., p. 465.
22. Ibid., p. 470.
23. Ibid.
24. Edward F. McClennen, “The Rationality of Being Guided by Rules,” in Alfred R. Mele and Piers Rawling, eds., *The Oxford Handbook of Rationality* (Oxford: Oxford University Press, 2004), esp. p. 234.
25. Ibid., p. 232.
26. Ibid., p. 238, n. 29.
27. Ibid., p. 224.
28. Ibid., p. 229.
29. See Edward F. McClennen *Rationality and Dynamic Choice: Foundational Explorations* (Cambridge, England: Cambridge University Press, 1990) and “Pragmatic Rationality and Rules,” *Philosophy and Public Affairs* 26 (1997).
30. David Gauthier, *Morals by Agreement* (Oxford: Oxford University Press, 1986).
31. Ibid., p. 186.
32. See David Gauthier “Assure and Threaten,” *Ethics* 104 (1994), and “Commitment and Choice: An Essay on the Rationality of Plans,” in Francisco Farina, Frank Hahn, and Stephano Vannucci, eds., *Ethics, Rationality, and Economic Behaviour* (Oxford: Oxford University Press, 1996).
33. Christopher McMahan, *Collective Rationality and Collective Reasoning* (Cambridge, England: Cambridge University Press, 2001).
34. Ibid., pp. 21–22.
35. Goldman, op. cit., p. 456 and p. 462.
36. I would like to thank Donald Bruckner and Dale Miller for commenting on the first draft and a later draft of this article, and Chrisoula Andreou for commenting on the first draft of it. Additionally, I would like to thank an anonymous referee for exceptionally helpful comments on this article, and Thomas Magnell, the Editor-in-Chief of the *Journal of Value Inquiry*, for his assistance.