

## *Conflicts of Rules in Hooker's Rule-Consequentialism*

BEN EGGLESTON<sup>1</sup>  
The University of Kansas  
Lawrence, KS 66045  
USA

It is the business of ethics to tell us what are our duties.

—John Stuart Mill<sup>2</sup>

### **I Introduction**

Just about any proponent of a rule-based theory of morality must eventually confront the question of how to resolve conflicts among the rules that the theory endorses. Is there a priority rule specifying which rules must yield to which, as in Rawls's lexical ordering of the first principle of his theory of justice over the second?<sup>3</sup> Must the agent intuitively bal-

---

1 I would like to thank Dale Miller for commenting extensively on several drafts of this paper, and two anonymous referees for this journal for their comments on the penultimate draft. I would especially like to thank Brad Hooker for his gracious encouragement of this project, and for his generous advice on it.

2 John Stuart Mill, *Utilitarianism* (in John Stuart Mill, *Essays on Ethics, Religion and Society* [volume 10 of the *Collected Works of John Stuart Mill*], J. M. Robson, ed. [Toronto: University of Toronto Press 1969], 203–59), at 219.

3 John Rawls, *A Theory of Justice*, revised ed. (Cambridge, MA: Harvard University Press 1999), 37–9.

ance considerations, as in certain forms of intuitionist pluralism? Or might there be some other conflict-resolving provision? Brad Hooker, a defender of a rule-based theory of morality that he calls ‘rule-consequentialism,’ confronts this question about conflicts of rules in his recent book *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*.<sup>4</sup> In this paper, I examine Hooker’s answer to this question, and I argue that his answer fails to solve a serious problem that arises from such conflicts.

In order to reach this conclusion, I’ll spend section II describing Hooker’s theory of morality, his answer to the question of how to resolve conflicts of rules, and the main problem that arises from such conflicts. In section III, I’ll argue that this problem is not solved, even in part, by Hooker’s answer. I’ll devote section IV to considering two objections that might be offered against my argument, before concluding in section V.

## II Hooker’s Rule-Consequentialism and Conflicts of Rules

Hooker presents his theory’s core proposition as follows. He writes,

There are many versions of rule-consequentialism. The version I favour is as follows:

RULE-CONSEQUENTIALISM. An act is wrong if and only if it is forbidden by the code of rules whose internalization by the overwhelming majority of everyone everywhere in each new generation has maximum expected value in terms of well-being (with some priority for the worst off). The calculation of a code’s expected value includes all costs of getting the code internalized. If in terms of expected value two or more codes are better than the rest but equal to one another, the one closest to conventional morality determines what acts are wrong. (32; see also 144, n. 3)

Having thus formally presented his theory’s standard of right and wrong, Hooker goes on to offer several clarifications and replies to objections.

‘One common misconception about rule-consequentialism,’ Hooker writes, has to do with how ‘rule-consequentialism [should] deal with conflicts between rules’ — that is, conflicts between two (or more) rules each of which seems to deserve a place within the ideal code (88-9). For example, both a rule requiring one to keep one’s promises and a rule requiring one to shield others from hardships of certain magnitudes would probably be included in the ideal code, but conflicts between

---

4 Brad Hooker, *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (New York: Oxford University Press 2000). Parenthetical references are to this work.

them can obviously arise (such as in a case in which a person can save another from considerable inconvenience only by breaking a promise to meet a third person at that same time). Hooker notes that some people think that rule-consequentialism should deal with such conflicts by 'build[ing] exceptions into the rules so as to keep them from conflicting' (89). And Hooker allows that the rules in the ideal code will have some exceptions built into them: 'For example, the rule about promise-keeping could have built into it an exception such that no one is required to keep a promise made to anyone who obtained the promise by lying' (89). But, he maintains, the exceptions built into the rules in the ideal code will not be sufficient to avert *all* conflicts between rules. The reason is not, of course, that there is anything inherently desirable about having some unresolved conflicts between rules, but because the exception clauses needed to avert all conflicts in acceptable ways would have to be so elaborate that a code whose rules were outfitted with such clauses would be nearly impossible for agents to learn, retain, and teach (89-90). To recall a phrase of Hooker's quoted above, the 'costs of getting the code internalized' (32) would be too high; 'At some point, the added costs involved in learning more complicated rules will outweigh the benefits' (90).

So Hooker faces the question of how conflicts of rules are to be resolved, and he rejects one of the obvious answers to this question. But why must he have an answer at all? What problem is caused by conflicts of rules, to which an answer to the question of conflicts of rules might be expected to provide a solution? The problem with which this paper is concerned might be called a problem of indeterminacy. When a theory's rules conflict in the way suggested above, the theory gives the agent no guidance as to which of her options she ought to choose. The agent remains torn between, say, saving the second person from inconvenience and keeping her promise to the third person. And surely this indeterminacy is a problem, since one of the main purposes of any moral theory is to yield judgments regarding which of the options an agent has in a given situation is the one she ought to choose. Obviously this is just a preliminary account of the problem; I'll return at the end of this section to the question of how serious this problem is.

In order to examine Hooker's proposal for dealing with conflicts of rules within his rule-consequentialist theory, it will be helpful to have a more precisely structured example in mind. Suppose that Mike is the president of a construction company that does a lot of road work for the city. For Mike this is not only fortunate but vital, since his high salary enables him to hire expert in-home supervision and productive interaction for his young son, who has a severe (but, with this labor-intensive treatment, gradually improving) psychological disorder that would otherwise require him to be institutionalized in a state facility where

his condition would only deteriorate, due to understaffing and a resulting reliance on managing patients through sedation. Mike believes that his company's steady stream of city contracts is due to his company's competitive bids and competent work, but in reality, his vice president for bidding and negotiations has found that the only way to get the city's supervisor of roads to take a company seriously is with bribery: a few thousand a month, in a brown paper bag slid casually across the table over lunch at a local deli. The vice president has kept this from Mike, knowing that the fewer people at the company who know about this arrangement, the better. The only person he's told is a trusted subordinate who helps him get the cash together every month.

One day, the vice president dies in a hunting accident, and Mike temporarily assumes his duties. The subordinate tells Mike about the arrangement, which Mike immediately regards as deeply immoral. But although he is loathe to let the current practice continue, Mike realizes that if he stops it, his company will stop winning contracts from the city and he and the company's other employees will lose their jobs, preventing him from continuing to keep his son's condition from worsening. Nor would publicly exposing the city's practices keep his company in business, since it would have to pay such massive fines as punishment for its many acts of bribery that any bankruptcy court would simply order it liquidated rather than allow it to preserve any of its capital. Finally, under no scenario would Mike have any hope of getting another decent job, since he would be seen as having presided over a company's suspicious demise or, if he goes public, as having been incredibly obtuse in having been unaware of such corruption going on right under his nose.

So Mike is forced to choose an act of one of these two types: (1) allowing a practice of bribery to continue and (2) allowing one's child to suffer a serious harm. For our purposes, two facts about acts of these two types are important. First, acts of each of these two types normally have bad consequences. Second, agents are quite infrequently confronted with circumstances in which they have no options aside from acts of these two types. Now, because of these facts, it is likely that the ideal code prohibits acts of each type *without making an exception permitting agents to avoid acts of the other type*, instead of containing a clause licensing an agent (in some or all cases of conflict) to perform an act of one type in order to avoid performing an act of the other type. For the more infrequently agents are confronted with circumstances in which they have no other options, the less work there is to be done by such a conflict-resolving clause, and thus the smaller are the benefits of a code's containing such a clause. Let us suppose, then, that in our hypothetical case the ideal code prohibits acts of allowing a practice of bribery to continue and acts of allowing one's child to suffer a serious

harm without saying that acts of one type are worse than acts of the other type.<sup>5</sup>

Now, what does rule-consequentialism have to say about such a case? Remember its standard of right and wrong — that an act is wrong if and only if it is forbidden by the ideal code of rules. In the case under consideration, each of Mike's options is forbidden by the ideal code of rules: the ideal code does not, we have stipulated, license Mike to choose one option in order to avoid the other. Of course in practice Mike would have to choose one option or the other (if only through inaction), but this does not affect the verdict of rule-consequentialism: according to it, neither of Mike's options is permissible; and neither is preferable to the other. So rule-consequentialism, as we understand it so far, is beset with the indeterminacy problem.

Hooker clearly regards conflicts of rules to be problematic and worth resolving. He responds by drawing on Richard Brandt's idea of the ideal conscience, where a conscience is understood basically as a set of aversions to performing acts of various kinds. And just as one code of rules may be better than another (in terms of the consequences of its internalization) in virtue of, say, requiring more (than the other code requires) to be done to alleviate poverty, so may one conscience be better than another in virtue of including stronger aversions (than the other conscience includes) to allowing poverty to persist without being alleviated. The ideal conscience, then, is ideal in just the way that the ideal code of rules is: it is that conscience whose internalization has maximum expected value (90). With this notion developed, Hooker writes the following:

If an act would be prohibited by the moral aversions thus selected [the moral aversions constitutive of the ideal conscience], it would be morally wrong, according to rule-consequentialism. When rules conflict, so do the aversions that are attached to them. The stronger aversion determines what action is permissible, according to rule-consequentialism. (90)

Then, following Brandt in regarding a *morally well-trained person* as a person who has the ideal conscience, Hooker approvingly quotes

---

5 Some might object, on grounds of public policy, that the consequences will obviously be better if one is required to eradicate corruption of the kind found in this case than if just one child, whose condition is not even life-threatening, continues to be the beneficiary of such expensive private treatment. Others, impressed by consequentialist arguments for agents' being specially attentive to those who are 'nearest and dearest,' might object that the consequences will obviously be better if one is required to give priority to one's family member. To such objectors, I say: please just worsen one of Mike's options, as needed, in order to make the two acts comparable in the way assumed in the argument.

Brandt's verdict that when rules conflict, the right thing to do is 'whatever course of action would leave morally well-trained people least dissatisfied' (90 and 132).<sup>6</sup>

The relevance of this dictum to the foregoing example is clear: a morally well-trained person would, like any agent, do *something* in the situation in question (even if only through inaction); and anything a morally well-trained person would do is something any agent is *allowed* to do. I take this employment, by Hooker, of the notions of the ideal conscience and of a morally well-trained person to be Hooker's answer to the question of conflicts of rules. In the next section, I shall explore this answer more fully, ultimately arguing that it fails to solve the indeterminacy problem.

But before closing this section, a few further points are worth noting. First, it may be observed — possibly with puzzlement — that rule-consequentialism as Hooker characterizes it on 90 is not consistent with rule-consequentialism as Hooker characterizes it on 32. For rule-consequentialism as characterized on 32 implies that an agent in the situation described above acts wrongly regardless of what she does, and rule-consequentialism as characterized on 90 avoids implying this. So what is supposed to be the relationship between these two passages? Are they supposed to be consistent? Is the apparently canonical statement of rule-consequentialism on 32 supposed to constrain all of Hooker's subsequent interpretive remarks? Or might some of Hooker's subsequent interpretive remarks function as revisions of what Hooker says on 32? I think that any reader of the book would have to agree that the last of these possibilities is the case. For it seems clear that what Hooker says on 90 (and on 131, where he notes that his theory can be understood in terms of *motivations*, not just rules) is meant to override what Hooker says on 32. Indeed Hooker notes as early as 2, and again on 91–2, that the notion of a conscience — not just the notion of rules — is essential to his theory. So any interpretive puzzlement should easily be dispelled.

Second, it might be wondered whether the indeterminacy in Hooker's theory is really a problem — or whether, on the contrary, regarding it as a problem depends on assumptions that a moral theorist may perfectly reasonably reject. If the latter is the case, then perhaps we should regard it as a pseudo-problem that Hooker should not be held responsible for solving.

---

6 Quoting Richard B. Brandt, 'Toward a Credible Form of Utilitarianism' (in *Morality and the Language of Conduct*, Hector-Neri Castañeda and George Nakhnikian, eds. [Detroit, MI: Wayne State University Press 1963], 107–43), at 134.

The main consideration in favor of regarding the indeterminacy problem as a (genuine) problem is simple and familiar. The essential mission of a moral theory, it is widely thought, is to articulate general principles that, in conjunction with specifications of agents' circumstances and options, yield judgments regarding which option an agent, in any given situation, ought to choose. This is not to say, of course, that the essential mission of a moral theory is to serve as a decision procedure that agents would be well-advised to employ in the course of everyday conduct. For it may be wise for agents to employ, as their decision procedures, simple rules of thumb that, as it turns out, do not amount to much of a theory. The best practice may not flow directly from the best theory, or from much of a theory at all.<sup>7</sup> But moral theories, even though their impact on practice may be indirect, are supposed to yield determinate verdicts as to which option an agent, in any given situation, ought to choose. Any indeterminacy is a deficiency that, other things being equal, it is desirable to eliminate.

There is, of course, plenty of support for this view in the tradition of moral-theory construction. Much of it is only implicit in the apparent aims of moral theorists, but some of them offer explicit advocacy of it. Mill, for example, does so in the epigraph to this paper. Moreover, Hooker himself shows some sympathy for this view. When surveying criteria for assessing moral theories, he asserts the methodological principle that 'Moral theories should help us to deal with moral questions about which we are not confident, or do not agree' (4). In defense of this principle, he writes that not only might a moral theory give us guidance in cases about which we are unsure,

it *should* do so (at least sometimes). For we turn to moral theory not merely to enlarge our understanding but to guide our practice... A moral theory that did not help with unsettled moral questions would thus let us down. (22)

In addition, after presenting his answer to the question of how conflicts of rules are to be resolved (the answer that we saw earlier in this section), he pointedly adds that 'rule-consequentialism is *not* crippled by conflicts between rules. It has a method for determining what is right in such situations' (91). Finally, Hooker writes that 'vagueness in rule-

---

7 On the distinction between moral theories (on the one hand) and decision procedures to be employed in practice (on the other), see Eugene Bales, 'Act-Utilitarianism: Account of Right-Making Characteristic or Decision-Making Procedure?' *American Philosophical Quarterly* 8 (1971) 257–65, and Peter Railton, 'Alienation, Consequentialism, and the Demands of Morality,' *Philosophy and Public Affairs* 13 (1984) 134–71.

consequentialism's conflict-resolving rules limits the theory's ability to revolve our moral uncertainties,' and concedes that 'Here is a place where rule-consequentialism comes up short' (136). So there is strong reason, both from the tradition of moral-theory construction and from within Hooker's own methodological framework, to regard indeterminacy as a problem worth solving. Admittedly, Hooker's views on indeterminacy are more complicated than I can explore here (see, for example, 127–36); all I claim is that the problem of indeterminacy is one to which Hooker offers a solution whose adequacy it is perfectly appropriate to assess.

### III An Assessment of Hooker's Solution to the Indeterminacy Problem

We saw above that Hooker's answer to the question of conflicts of rules consists, essentially, of his employment of the notions of the ideal conscience and of a morally well-trained person. And it might appear that his employment of these notions does, indeed, solve the indeterminacy problem. For with these notions in play, Hooker is equipped to say that according to rule-consequentialism, the right thing to do is 'whatever course of action would leave morally well-trained people least dissatisfied' (quoted above, in section II). And this might appear to provide just the sort of guidance that is sufficient to address the indeterminacy problem. In my view, however, Hooker's employment of these notions offers no solution to this problem.

Before proceeding, I should underscore the strength of my thesis by distinguishing it from another, weaker, thesis. The weaker thesis from which I want to distinguish mine is the claim that Hooker's employment of these notions does not *completely* solve the indeterminacy problem. But this thesis is trivial because Hooker does not suggest otherwise; on the contrary, there are at least three distinct passages — all subsequent to Hooker's employment of the notions of the ideal conscience and of a morally well-trained person — in which (on my reading of them) Hooker says that some ineliminable indeterminacy remains in his theory.<sup>8</sup> So this weaker thesis is not at issue here. My thesis, rather,

---

8 On 107, he writes, 'Rule-consequentialist agents will have to rely on judgement to resolve conflicts between general duties.' On 116, he writes that in some cases, 'rule-consequentialism mirrors the indeterminateness of conventional morality.' Finally, on 133, he writes, 'Judgement will ... be needed to resolve conflicts where the degree of one consideration at stake is less dramatically different from the degree of the other consideration at stake.'



is that Hooker's employment of these notions does not solve the indeterminacy problem *at all*. And here Hooker and I are in disagreement. For the very question that Hooker intends to answer with these notions is the question of conflicts of rules. Now it might be wondered whether, in employing these notions, Hooker aims to solve some other problem than the indeterminacy problem and is happy to leave the indeterminacy problem entirely unsolved. But this seems unlikely. For in discussing conflicts of rules, Hooker does not give significant attention to any other problem that arises from such conflicts. And as we saw in the last paragraph of the previous section, Hooker claims that 'rule-consequentialism is *not* crippled by conflicts between rules. It has a method for determining what is right in such situations' (91). Presumably, in referring to determining what is right, Hooker means to be referring to the lessening of indeterminacy; and the method to which he refers is, of course, the one based on the notions of the ideal conscience and of a morally well-trained person. So Hooker clearly regards his employment of these notions as solving the indeterminacy problem *at least to some extent*.

As I said, I deny this claim. For Hooker's employment of these notions would make a dent in the indeterminacy problem only if it were possible for the ideal conscience to be *more* determinate than the ideal code of rules or, in other words, for (1) the ideal conscience to contain a stronger aversion to acts of one type than to acts of another type, while (2) the ideal code of rules did *not* indicate that acts of one type are worse than acts of the other type. To see what this state of affairs would amount to in a particular case, recall the example from above, in which Mike must choose between an act of allowing a practice of bribery to continue and an act of allowing one's child to suffer a serious harm. The ideal rules, it will be recalled, are silent on this choice. Now in order for the ideal conscience to be more determinate than the ideal rules, it would have to be possible for the ideal conscience to contain a stronger aversion to acts of one of these types than the other. (I do not claim that the ideal conscience would have to *actually* contain a stronger aversion to acts of one of these types than the other, since Hooker's employment of the notions of the ideal conscience and of a morally well-trained person would deliver *some* improvement in determinacy even if it didn't provide determinacy in every case, as long as it provided determinacy in some cases. So I claim only that, in order for Hooker's employment of these notions to succeed in solving the indeterminacy problem even in part, it would have to be *possible* for the ideal conscience to contain a stronger aversion to acts of one of these types than the other.)

It might seem that, even while the ideal code of rules does not imply that acts of one type are worse than acts of the other type, the ideal conscience could well contain a stronger aversion to acts of one type

than to acts of the other type. For how determinate a code of rules or a conscience is depends heavily on how complex it is, and it might be assumed that, for some given level of complexity, the internalization costs of a code of rules of that level of complexity would typically be higher than the internalization costs of a conscience of that level of complexity. If this assumption — to which I'll turn in the next paragraph — were true, then one might expect that the ideal conscience would be more complex, and hence more determinate, than the ideal code of rules. Indeed, this would be quite reasonable to expect, given the aforementioned assumption about differential internalization costs, in the same way that, when someone is buying a laptop computer and a desktop computer, it is quite reasonable to expect that the desktop computer will contain more memory than the laptop computer because memory is cheaper for desktop computers than for laptop computers. So the assumption of lower internalization costs for consciences than for codes of rules leads quite naturally to the thought that the ideal conscience would be more determinate than the ideal code of rules.

Let us turn, then, to the assumption of lower internalization costs for consciences than for codes of rules. To be sure, this assumption has some intuitive plausibility. For it might be thought that internalizing a code of rules would require explicitly cognizing and memorizing a lengthy series of propositions, and that one would not have truly internalized a code unless one were capable of reciting it on command. (Of course being capable of that could not reasonably be thought sufficient for having internalized that code; obviously some people are capable of reciting codes they repudiate. But it might be thought that someone who cannot recite a code of rules has not really internalized it.) On the other hand, developing the corresponding conscience (i.e., the conscience with the same content) might look easier, since it just involves acquiring certain dispositions, such as dispositions to perform acts of certain types, dispositions to refrain from acts of certain other types, and (usually) dispositions to counsel and pressure others to behave similarly. Acquiring these dispositions is surely easier than memorizing a long list of rules, since (as reflection on the way people are actually brought up suggests) it does not involve a lot of conscious memorization, and does not involve acquiring the ability to recite much, if anything, on command.

It turns out, however, that on Hooker's account, internalizing a particular code of rules is *no more costly* than developing the corresponding conscience, since — again, on Hooker's account — to internalize a particular code of rules *just is* to develop the corresponding conscience. That is, what Hooker means by internalizing a particular code of rules is not something separate from, or above and beyond, or experientially distinct from, developing the corresponding conscience. Rather, de-

veloping the corresponding conscience is all there is to internalizing a particular code of rules. Indeed in the opening paragraphs of his book, Hooker remarks that 'The collective internalization of a code [of rules] amounts to the establishment of a shared conscience' (2). And this is not, on Hooker's part, just a stipulative definition of internalizing a code of rules, meant to make this term synonymous with developing the corresponding conscience. Rather, this claim is substantiated by the account of internalizing a code of rules that Hooker offers in his book. For as Hooker's account proceeds, it becomes apparent that internalizing a particular code of rules is in some ways easier, but is in other ways more transforming, than acquiring the ability to recite a list of rules on command.

Internalizing a particular code of rules is easier than learning to recite it, Hooker writes, because "We can learn things without being able to recite them" (96). To explain this point, Hooker approvingly quotes Brandt as saying that what internalizing a particular code of rules involves is knowing it

well enough to recall the relevant rule when stimulated by being in a context to which it is relevant. Learning a moral code is thus like learning a complex route into a large city: we may not be able to draw it or explain to others what it is, but when we drive it and have the landmarks before us, we remember each turn we are to make (97).<sup>9</sup>

Alas, the analogy is not entirely apt, since driving into a large city via a certain route presents the agent with essentially the same stimuli every time, and presumably learning a moral code involves being able to react appropriately to new stimuli or to old stimuli arranged in new ways. So perhaps learning a moral code is more like learning one's way around a large city (as a taxi driver might) than learning a particular route into it. But Brandt's point is clear enough. Hooker also approvingly quotes David Copp:

[I]t is possible to intend to conform to a rule, or to desire its currency, even if one does not know how to formulate it. One must be *aware* of the rule, but this does not require knowing its formulation. However, if a person subscribes to a moral standard, she must be able to recognize conformity with the standard in a wide range of circumstances and to see what would count as conformity. I do not mean that her recognition must be infallible and without difficulty or doubt; I mean only that she can recognize what counts as conformity with rough accuracy. (97)<sup>10</sup>

---

9 Quoting Brandt, 'Towards a Credible Form of Utilitarianism,' 142, n. 6.

10 Quoting David Copp, *Morality, Normativity, and Society* (New York: Oxford University Press 1995), 86.

So internalizing a particular code of rules does not require learning to recite it, and thus in some ways internalizing a particular code of rules is easier than learning to recite it.

But in other ways, learning a particular code of rules is more transforming, for the agent, than learning to recite it. For according to Hooker, accepting a certain code of rules

incorporates not just the disposition to comply with these rules. Acceptance of rules also involves the disposition to encourage others to comply with them, dispositions to form favourable attitudes towards others who comply with them, dispositions to feel guilt and shame when one breaks them and to condemn and resent others' breaking them, all of which dispositions and attitudes being supported by a belief that they are justified. (76)

Hooker subsequently suggests that having internalized a particular code of rules is equivalent to having developed a conscience of a certain kind (82-3), and in the following passage he completes the connection between codes of rules and consciences:

Rule-consequentialism takes the acceptance of rules to involve *more than* certain associated motivations. It also involves having sensitivities, emotions, and beliefs — indeed a particular cast of character and conscience. If you *accept* a rule against stealing, you will be motivated not to steal simply because it is stealing (not merely because you will get into trouble). You will also be disposed to feel guilty if you steal, disposed to resent stealing by other people, and disposed to blame them for it. You will want others to have these dispositions not to steal and to react negatively to those who do steal. And you will have associated beliefs, such as that stealing is morally prohibited and that this prohibition is justified. We might sum all this up by saying that to accept a code of rules is just to have *a moral conscience of a certain shape*. In other words, when rule-consequentialists consider alternative codes of rules, they are considering alternative possible contours for people's consciences. (91)

It should be clear, then, that internalizing a particular code of rules is in some ways easier, but is in other ways more transforming, than acquiring the ability to recite a list of rules on command. More to the point, it should be clear that Hooker regards internalizing a particular code of rules as the same thing — the same activity, experience, or event — as developing the corresponding conscience.

To see the significance of this equivalence, let us review the main ideas of this section. Our principal concern is whether Hooker's employment of the notions of the ideal conscience and of a morally well-trained person solves, even if only partially, rule-consequentialism's indeterminacy problem. I claimed that Hooker's employment of these notions accomplishes this only if it is possible for the ideal conscience to be more determinate than the ideal code of rules. I then explored an argument that might seem to support this possibility — an argu-

ment based on the assumption of lower internalization costs for consciences than for codes of rules. With these considerations in view, we can see that the foregoing equivalence (between internalizing codes of rules and developing consciences) has a double significance. First, it rebuts the foregoing argument in support of the possibility of the ideal conscience's being more determinate than the ideal code of rules, by undercutting the assumption of lower internalization costs for consciences than for codes of rules — for if internalizing a code of rules and developing the corresponding conscience are one and the same, then it doesn't make sense even to distinguish them as events, much less to regard one event as less costly than the other.

Second, the foregoing equivalence provides the lynchpin of a straightforward affirmative argument showing that Hooker's employment of the notions of the ideal conscience and of a morally well-trained person cannot solve, even partially, rule-consequentialism's indeterminacy problem. For if internalizing a particular code of rules is equivalent to developing the corresponding conscience, then the ideal code of rules has the same content as the ideal conscience, and the ideal conscience is no more complex than the ideal code of rules. From this it follows, of course, that the ideal conscience cannot be more determinate than the ideal code of rules. I conclude, then, that Hooker's employment of the notions of the ideal conscience and of a morally well-trained person does not reduce, much less eliminate, the indeterminacy of rule-consequentialism.

#### IV Two Objections

In this section, I consider two objections to the concluding claim of the previous section. One objection highlights the actions that morally well-trained people actually choose to perform; the other draws on the potential complexity of the ideal conscience.

First, I claimed above that Hooker's employment of the notions of the ideal conscience and of a morally well-trained person does not reduce the indeterminacy of rule-consequentialism. This claim may appear to be refuted by Hooker's claim that when rules conflict, the right thing to do is 'whatever course of action would leave morally well-trained people least dissatisfied' (quoted above, in section II). For surely, when a morally well-trained person is faced with a choice of the kind with which we are concerned (such as the choice Mike faces), she will do *something*, because of being least dissatisfied with it. And if (as Hooker says) rule-consequentialism says that *this* is the right thing to do, then however plausible the argument of section III may appear, its conclusion is plainly false: Hooker's employment of the notions of the ideal

conscience and of a morally well-trained person *do* make rule-consequentialism more determinate.

In response to this line of reasoning, I grant that in any given situation, every morally well-trained person would do something (that she is least dissatisfied with). I just deny that in any given situation, there is something — one particular act — that every morally well-trained person would do (or, to be precise, that would leave every morally well-trained least dissatisfied). Of course, in some situations, there is only one particular act the performance of which is compatible with being a morally well-trained person. But in situations of conflicts of rules, this is not true. In such situations, there are multiple options any of which would leave a morally well-trained person least dissatisfied. In such situations, every morally well-trained person would do something, but only by drawing on other parts of her deliberative capabilities than those constitutive of her being a morally well-trained person — i.e., her conscience (which is the ideal conscience, since she is a morally well-trained person). For example, a person typically has self-interested preferences over her options, and even though those self-interested preferences may get their full moral due by way of her conscience (since any reasonable moral theory, including Hooker's, allows a person to give some weight to her own interests in making a moral decision), even the most morally conscientious person may not hesitate to bring these self-interested preferences to bear in order to resolve an indeterminacy that remains once moral considerations have had their say. Another, more basic, example is found in the fact that most people are not like Buridan's ass: they are perfectly capable of choosing among multiple, equally good, options and getting on with their lives, such as when a shopper manages to buy cereal in the supermarket without getting hung up on which box of her favorite cereal to put in her box. So when a morally well-trained person does something, she does not necessarily do it *qua* morally well-trained person, period; she may well do it *qua* morally well-trained person who, in addition to being a morally well-trained person, has other deliberative resources that are capable of filling in the gaps left by her conscience's indeterminacy.

I have been arguing that (1) the fact that rule-consequentialism says that the right thing to do is whatever a morally well-trained person would do and (2) the fact that a morally well-trained person would do *something* do not combine to imply that (3) the first fact makes rule-consequentialism more determinate than it would otherwise be. To see this point by way of a brief analogy, suppose that I were to propound the following principle of right and wrong: an act is right if and only if its consequences are better than the consequences of *some* act open to the agent in that situation. That is, all an agent has to do to act rightly in a given situation is to avoid doing the worst available act (or one of

the worst available acts, if there is a tie for worst), where badness is understood in terms of the consequences of the particular acts available to the agent. (So this is a non-maximizing — indeed an extremely modest and undemanding — variant of act consequentialism.) This principle is vulnerable to many objections, not the least of which is that it is woefully indeterminate. For an agent faced with a choice typically has more than two options open to her, and just ruling out the *worst* option as what she should *not* choose leaves entirely unspecified which of her remaining options is the one that she ought to choose. Now suppose that, in order to respond to this objection, I invoke the notion of a morally well-trained person, where this notion is defined as Hooker defines it (adjusted, of course, to reflect the different content of the moral principle presently under discussion). Then suppose that, having introduced the notion of a morally well-trained person, I purport to solve (at least in part) the indeterminacy problem by saying that according to my theory, the right thing to do is whatever a morally well-trained person would do. For surely (my supposed solution goes) a morally well-trained person would do *something*; and thus surely it mitigates my theory's indeterminacy problem if I say that whatever a morally well-trained person would do is the right thing to do. But of course this is no solution. Although a morally well-trained person would do *something*, I do not make any dent in my theory's indeterminacy problem by saying that the right thing to do is whatever a morally well-trained person would do. And neither, I submit, does Hooker.

The second objection I want to consider draws on the potential complexity of the ideal conscience in order to counter my claim that Hooker's employment of the notions of the ideal conscience and of a morally well-trained person does not reduce the indeterminacy of rule-consequentialism. This objection maintains that even if it is established that internalizing a particular code of rules is the same event or experience for a person as developing a conscience with a certain content, it does not follow that the content of the conscience must be limited to that of the code in question. That is, while the content of the ideal code of rules might run out at a certain point, the content of the ideal conscience might extend further, making the ideal conscience more determinate than the ideal code. To take a simple example, the ideal code might contain a rule against lying and a rule requiring one to promote one's family members' important interests, without specifying which rule should take precedence in cases of conflict. Nevertheless, in the ideal conscience, the aversions corresponding to these rules — the aversion to lying and the aversion to letting one's family members down—might be of different strengths in a certain subset G of the class of cases in which these rules come into conflict, so that the latter aversion is stronger than the former in such cases. For example, suppose Sandy has the

ideal conscience, and is faced with a situation in which she has no way to send her daughter to college except by lying on a tax form — i.e., she must choose between lying on a tax form and failing to send her daughter to college. Suppose also that this case falls within subset G. Then it follows from our suppositions about Sandy and the ideal conscience that Sandy's aversion to letting her daughter down would be stronger than her aversion to lying on the tax form. Of course, which aversion would be stronger, in the ideal conscience, is an empirical question; maybe the ideal conscience would handle cases in subset G the other way around. But this contingency does not affect the point of the objection: that there is no reason why one aversion could not be stronger than the other in cases in subset G, thus making Sandy's conscience — the ideal conscience — more determinate than the ideal code.<sup>11</sup>

This objection turns on a subtlety in the notion of the ideal conscience that merits further attention, but that I do not think ultimately succeeds in solving the indeterminacy problem. First, though, in order to see how this example gives rise to an objection to my claim, it is important to distinguish it from two related but unilluminating ones for which it may easily be mistaken. The first related example is one in which Sandy is, indeed, disposed to handle cases in subset G by lying rather than letting family members down, but only due to deliberative resources of hers that fill some of the gaps left by her conscience — deliberative resources such as a desire to help her family members or an inclination to avoid harms to identifiable individuals more assiduously than “impersonal” harms such as that of tax evasion. As I mentioned earlier in this section (in my discussion of the deliberative resources that we can cite in order to explain how a morally well-trained person chooses something to do in a case of indeterminacy), such preferences and motivations may get their full moral due by way of the ideal conscience, but this does not mean that an agent — even one with the ideal conscience — will not bring them to bear once moral considerations (her conscience) have had their say and left her in a situation of indeterminacy. Obviously, if such extra deliberative resources are the explanation for Sandy's being disposed to handle cases in subset G by lying rather than letting her family members down, then we have no reason to credit her conscience with the extra determinacy needed to account for her conduct in cases in subset G. So the Sandy example gives rise to an objection to my claim only if we construe it so that it is Sandy's *conscience* (and not other deliberative resources of hers) that leads her to handle cases in subset G in a certain way.

---

11 I owe this objection to an anonymous referee for this journal.



The second related but unilluminating example with which we must not confuse the useful Sandy example is one in which Sandy was given an education meant to instill in her the ideal conscience, but this resulted in Sandy's having a conscience more determinate than the ideal one. This might have happened in any number of ways. One possibility is that Sandy and others in her generation were (in their families, schools, clubs, community associations, etc.) given the sort of moral education that, on average, would result in their having the ideal conscience, but in Sandy's case it resulted in a conscience more determinate than the ideal one, perhaps due to Sandy's above-average aptitude for grasping distinctions or her being more inclined than most people to find moral principles covering all conceivable cases. Her conscience may even be *better* than the ideal one, in the sense that — setting internalization costs aside — the widespread internalization of Sandy's conscience would have better expected consequences than the widespread internalization of the ideal conscience. (In such a case the most likely explanation for Sandy's conscience being less than ideal would be that the widespread internalization of Sandy's conscience would be much more costly to achieve and sustain, resulting in a lower *net* expected value for its widespread internalization than for that of the ideal one.) On the other hand, her conscience, though more determinate than the ideal one, might have worse expected consequences, perhaps due to Sandy's conscience containing some distinctions and directives that do more harm than good. Either way, this example is one in which Sandy's conscience is not the ideal one; it is more determinate than the ideal one. So this example does not show how the ideal conscience can be more determinate than I have claimed it can be.

The lesson of the last two paragraphs is that, in order for the Sandy example to support the objection under consideration, it must be set up so that the way Sandy handles cases in subset G (i.e., by lying rather than her letting family members down) is due to her conscience, and it must be set up so that Sandy really just has the ideal conscience — not a more determinate one. And all this must be true even though the ideal code of rules does not dictate a particular way for agents to handle cases in subset G. It is doubtful, however, that such a scenario is permitted by Hooker's construal of the notions of a code of rules and the conscience that corresponds to it. Two lines of reasoning give rise to these doubts.

The first line of reasoning begins with the fact that Hooker repeatedly implies that, for each code of rules, there is one corresponding conscience. For example, as quoted in the previous section, Hooker writes that

to accept a code of rules is just to have a *moral conscience of a certain shape*. In other words, when rule-consequentialists consider alternative possible codes of rules, they are considering alternative possible contours for people's consciences. (91; emphasis in original)

We are entitled, then, to think of every code of rules  $R_i$  as having just one conscience — call it  $C(R_i)$  — that corresponds to it. And if considering  $C(R_i)$  is a way for rule-consequentialists to consider  $R_i$  (as Hooker says is the case), then the content of  $R_i$  and the content of  $C(R_i)$  must be related in some systematic way, for every  $i$ . Identity, obviously, is one systematic way for the two contents to be related: for every  $i$ , the content of  $R_i$  is the same as the content of  $C(R_i)$ , in the sense that what the rules of  $R_i$  require is the same as what the aversions of  $C(R_i)$  would lead a person to do. Are there other plausible candidates than identity? It will not work to say that the way the contents are related is that  $C(R_i)$  is systematically more determinate than  $R_i$ , since this way of relating  $R_i$  and  $C(R_i)$  does not shed any light on how any extra determinacy possessed by  $C(R_i)$  translates into particular content (e.g., whether agents with some  $C(R_i)$  would side with truth, not family, in cases in subset  $G$ , or vice versa). Nor will it work to characterize the relation in terms of anything like the consequences of widespread internalization, since we need to know the contents of  $R_i$  and  $C(R_i)$ , for any  $i$ , in order to identify and compare the consequences of their widespread internalization. It is hard to imagine, then, how considering consciences can be equivalent to considering codes of rules, as Hooker says is the case, unless the content of the conscience corresponding to some particular code of rules is simply the same as the content of that code of rules itself.

A second line of reasoning in support of this conclusion brings us back to the case of Sandy; in particular, it appeals to how we would be obliged to describe the code of rules she has internalized — call it  $R_{Sandy}$  — if we knew certain things about the aversions constitutive of her conscience. Suppose, for example, that we knew that Sandy's conscience contains a stronger aversion to letting her family down in cases in subset  $G$  than to lying in such cases. Then what should we say about  $R_{Sandy}$ ? Would it make sense to say that  $R_{Sandy}$  is indeterminate in such cases? Or should we say that  $R_{Sandy}$  has some determinacy, in that it contains a priority rule saying that the family-members rule overrides the rule against lying in cases in subset  $G$ ? To provide criteria for answering these questions, I would suggest that any adequate account of  $R_{Sandy}$  must do two things, above all else: (1) it must accurately reflect what it is like for Sandy when she engages in moral deliberation and (2) it must underwrite successful explanations and predictions of Sandy's behavior. If the aversions in Sandy's conscience have the strength differential hypothesized above, then one aspect of what moral deliberation is like,

for Sandy, is that she perceives the needs of her family as taking precedence over truthfulness in cases in subset G. Obviously we do not capture this aspect of Sandy's experience as a moral deliberator unless we credit  $R_{Sandy}$  with containing the sort of priority rule mentioned above. Similarly, if the aversions in Sandy's conscience have the hypothesized strength differential, and this is reflected in her handling of cases in subset G in a certain way, then our account of  $R_{Sandy}$  will be explanatorily and predictively inadequate unless, again, we credit it with containing the sort of priority rule mentioned above.

In response to this line of reasoning, there are four objections on which I want to comment, albeit briefly. First, it might be claimed that the two criteria of adequacy I mentioned above for our account of  $R_{Sandy}$  are peripheral or irrelevant, and that our account of  $R_{Sandy}$  should be governed by other considerations — ones, as it happens, that do not militate in favor of construing  $R_{Sandy}$  in the way I propose. I cannot briefly show this objection to be groundless, but I think that exploring it would take us too far into the obscure subject of criteria of adequacy for accounts of codes of rules internalized by individuals, so I shall have to rely on whatever intuitive appeal those criteria possess, and just acknowledge that the conclusion of this paper is conditional on the premise that those criteria have the primacy I claim for them.

Second, there are different ways of stating the specific fact about  $R_{Sandy}$  that I have been arguing we must acknowledge, and that I have been referring to in terms of a priority rule. For example, instead of saying that  $R_{Sandy}$  contains a priority rule of a certain kind, we might equally well say that in  $R_{Sandy}$ , the rule against lying has an exception clause making it yield to the family-members rule in cases in subset G. But the fact that none of the legitimate descriptions of this feature of  $R_{Sandy}$  is privileged, or canonical, does not mean that we are entitled to omit all of them, and to maintain that  $R_{Sandy}$  lacks the determinacy that all these formulations describe.

Third, it might be true, of course, that Sandy was never deliberately taught to handle cases in subset G (or some superset of G, etc.) by giving priority to her family members. But we can easily reconcile this with our preferred account of  $R_{Sandy}$  by pointing to the fact that people often internalize rules (including priority rules) that they were not deliberately taught. Finally, it might well be true that Sandy does not report the priority rule in question when asked to provide an account of her code of rules, and even denies the presence of it when asked to reflect on it specifically. Again, however, I would suggest that we should maintain our preferred account of  $R_{Sandy}$ , if her aversions really have the hypothesized strength differential; we should sooner fault Sandy for a small lack of self-knowledge than fail to meet the criteria of adequacy stated above.

I have been arguing, in answering the second objection introduced in this section, that any code of rules should be construed as having just the same content as the conscience corresponding to it. I have argued that this is required for two purposes: to make good on Hooker's claim that considering consciences is equivalent to considering codes of rules, and to have a phenomenologically adequate, and explanatorily and predictively adequate, account of the code of rules internalized by any particular individual. Surely this conclusion must hold for the *ideal* code of rules and the conscience corresponding to it (the ideal conscience), since the ideal code of rules, and the ideal conscience, do not differ from their non-ideal rivals in virtue of any features that exempt them from the foregoing considerations; they are just ranked higher than their rivals in a certain kind of cost-benefit calculation involving beneficial behavior, internalization costs, and so on. The content of the ideal conscience, then, cannot extend beyond that of the ideal code of rules, and Hooker's employment of the notions of the ideal conscience and of a morally well-trained person does not reduce the indeterminacy of rule-consequentialism.

## V Conclusion

In closing, I want to acknowledge one difference that Hooker's employment of the notions of the ideal conscience and of a morally well-trained person does make in the context of the indeterminacy problem. The difference is this: by equating internalizing a code of rules with developing the corresponding conscience, Hooker makes the extent of rule-consequentialism's indeterminacy less than one might think it would be, if one were to assume that internalizing a code of rules requires acquiring the ability to recite a list of rules on command. For if internalizing a code of rules were like that, then the internalization costs for a given code of rules would probably be higher than they actually are, on Hooker's account, and the ideal code of rules would be less complex, and hence more indeterminate, than it actually is, on Hooker's account. So the extent of the indeterminacy in Hooker's rule-consequentialism is less than in some other conceivable rule-based consequentialist theories.

My concern, though, is that indeterminacy remains, and the notions of the ideal conscience and of a morally well-trained person cannot be employed to reduce it. That is, if we are considering an agent's situation from a third-person perspective (e.g., trying to advise her or evaluate her conduct), and we know what the ideal code of rules says, then if this code of rules is indeterminate as to how the agent should act, the notions of the ideal conscience and of a morally well-trained person have

nothing to offer us: their indeterminacy-reducing powers have been entirely used up in the additional complexities that they enable the ideal code of rules to contain. Similarly, if an agent who has internalized the ideal code of rules faces a choice on which that code of rules is silent, she cannot hope to get any guidance by inquiring into what a morally well-trained person would do, because (by hypothesis) she *is* a morally well-trained person and doesn't know what to do. So the notions of the ideal conscience and of a morally well-trained person do not remove any indeterminacy from rule-consequentialism. Now it may seem at this point that my thesis is reducible to the complaint that once the benefits of the notions of the ideal conscience and of a morally well-trained person have been gained, then there is nothing to be gained from them. But trivial though this formulation may seem, it states something significant. For it is natural to read Hooker's answer to the question of conflicts of rules as suggesting that the notions of the ideal conscience and of a morally well-trained person *can* usefully be invoked in third-person and first-person contexts such as those just mentioned, in order to solve the indeterminacy problem. And yet, I argue, they cannot.

Let me now put my claims in a somewhat broader context. First, I believe it is perfectly reasonable for Hooker to liken a person's internalization of the ideal code of rules to her development of the ideal conscience and, one could say, to her becoming a morally well-trained person. Second, I believe it is perfectly correct to regard Hooker's employment of these notions as rendering his theory one that is more determinate than some other rule-based consequentialist theories are. All I claim is that it is *not* correct to regard Hooker's employment of these notions as solving, even in part, whatever indeterminacy problem rule-consequentialism in fact has.

I do not know how weighty my claims are. I do not think they are weighty enough, by themselves, to offset the many strengths of Hooker's very clear, original, sophisticated, and well-defended rule-consequentialist theory. On the other hand, I do think that they establish that the indeterminacy problem is unmitigated by Hooker's proposed solution to it. In any case, I think it can be concluded that conflicts of rules remain problematic for Hooker's rule-consequentialism.

*Received: June 2005*

*Revised: March 2006*