Solving Rule-Consequentialism's Acceptance Rate Problem

TIMOTHY D. MILLER

Lee University

Recent formulations of rule-consequentialism (RC) have attempted to select the ideal moral code based on realistic assumptions of imperfect acceptance. But this introduces further problems. What assumptions about acceptance would be realistic? And what criterion should we use to identify the ideal code? The solutions suggested in the recent literature – Fixed Rate RC, Variable Rate RC, Optimum Rate RC and Maximizing Expectation Rate RC – all calculate a code's value using formulas that *stipulate* some *uniform* rate(s) of acceptance. After pointing out a number of difficulties with these approaches, I introduce a formulation of RC on which *non-uniform* acceptance rates are *calculated* rather than *stipulated*. In addition to making more realistic assumptions about acceptance rates, *Calculated Rates RC* has several other advantages: it gives equal consideration to both *acceptance* and *compliance* rates and it brings RC more in line with our intuitive ways of thinking about rules and their consequences.

Rule-consequentialism (RC) offers an attractively simple conception of moral theory: the goal is to identify the *ideal moral code* – i.e. the complete set of moral rules that would have better consequences than any alternative set. But beyond the simplicity of this basic idea, RC is beset by complexities of formulation. Perhaps the most fundamental issue concerns specifying the criterion for evaluating the consequences of competing moral codes. The simplest criterion would select the code that would have the best consequences if it were complied with. However, in the real world, non-compliance is an unavoidable problem, and the simple criterion would not select a code that includes rules concerning how to deal with non-compliance. Furthermore, codes have consequences that extend beyond those of compliance; the *internalization* or *acceptance* of a code can influence attitudes, expectations and behaviours in ways that cannot be captured by simple compliance formulations. Focusing on acceptance also highlights the costs associated with teaching a rule – costs that are important to consider if the theory is to avoid selecting rules that are overly complicated or excessively demanding.[1] In light of such considerations, recent formulations of RC have attempted to compare the consequences of codes on the basis of realistic assumptions of *imperfect* or *partial acceptance*. But this introduces a further problem: what assumption about acceptance would be realistic?

[1] B. Hooker, *Ideal Code, Real World* (Oxford, 2000), pp. 75–80.

© Cambridge University Press 2015
Utilitas Vol. 28, No. 1, March 2016
doi:10.1017/S0953820815000229

CrossMark

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

All of the solutions to these problems proposed in the recent literature calculate a code's value using formulas that *stipulate* some *uniform* rate(s) of acceptance. After pointing out a number of difficulties with these approaches, I will introduce an alternative formulation of RC on which *non-uniform* acceptance rates are *calculated* rather than stipulated. In addition to making more realistic assumptions about acceptance rates, *Calculated Rates RC* has several other advantages over previous formulations: it gives equal consideration to both *acceptance* and *compliance* rates, and it brings RC more in line with our intuitive ways of thinking about rules and their consequences.

FOUR RECENT FORMULATIONS OF RC

Fixed Rate RC

Brad Hooker's nuanced defence of RC in *Ideal Code, Real World* is the starting point from which recent alternative formulations of RC have departed. Hooker characterizes the ideal code as the 'code of rules whose internalization by the overwhelming majority of everyone everywhere ... has maximum expected utility'.[2] Hooker later renders the vague characterization 'overwhelming majority' more precise by stipulating an acceptance rate of 90 per cent. Of course, this stipulation trades vagueness for a measure of arbitrariness, but Hooker defends the choice of 90 per cent by noting that this rate has the advantages of being high enough to support the idea that the rules are intended to be internalized and followed by everyone, while also including a rate of non-acceptance high enough to ensure adequate rules concerning how to handle non-acceptance and non-compliance.[3] So Hooker's Fixed Rate RC (FRRC) characterizes the ideal code as follows:

> (FRRC) An *ideal code* is a code whose expected value at a 90 per cent acceptance rate is at least as high as any other code.[4]

In spite of Hooker's defence of his chosen rate, the arbitrary nature of the stipulation has been a common point of criticism. Michael Ridge has objected that rules that would deal well with a 10 per cent rate of non-acceptance might be horribly ineffective in the face of, say, 20 or 30 per cent non-acceptance.[5] Ridge contends that an ideal code should include provisions to deal with all of the real world possibilities that

[2] Hooker, *Ideal Code*, p. 32.

[3] Hooker, *Ideal Code*, p. 84.

[4] If several codes are tied for having the highest expected value, Hooker gives preference to the code most similar to conventional morality. See Hooker, *Ideal Code*, pp. 114–17. For present purposes, I will ignore this complication.

[5] M. Ridge, 'Introducing Variable-Rate Rule Utilitarianism', *The Philosophical Quarterly* 56 (2006), pp. 242–53, at 244–6. Several of the theories considered in this article are formulated as versions of rule-utilitarianism rather than rule-consequentialism.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

might arise, and he argues that an approach that arbitrarily stipulates some fixed acceptance rate would be unable to do so.

Variable Rate RC

Ridge's solution is to evaluate codes by considering their expected values at all levels of acceptance with an eye towards identifying the code with the highest *average* expected value. Ridge refers to this theory as Variable Rate RC (VRRC):

> (VRRC) An *ideal code* is a code whose average expected value for all acceptance rates is at least as high as the average expected value of all alternative codes.[6]

Although VRRC does not arbitrarily select an acceptance rate, Kevin Tobia has argued that the theory remains arbitrary in another respect. There are a variety of ways that multiple acceptance rates could be factored into the comparative evaluation of codes: we could select the code with the highest average value (as VRRC does), but we could also select the code with the highest median, the highest mode, the highest minimum value, and so on. Tobia objects that 'it is arbitrary to choose one [of these criteria] without considering the others.'[7]

An even more fundamental problem for VRRC is the apparent irrelevance of its criterion. The charge of irrelevance has been made in a variety of specific ways,[8] but the general underlying problem is that VRRC's calculation gives equal weight to a code's value at every rate of acceptance. However, there are good reasons to deny that all acceptance levels are of equal significance. For example, Hooker insists that a moral code is, by its very nature, intended to apply to everyone, and it cannot fulfil this intended function without widespread acceptance.[9] If he is right about that, and I think he is, then it would seem to follow that relatively high acceptance rates should matter more than relatively low rates.

Another way to see just how counterintuitive it would be to grant equal significance to low acceptance rates is to note the obvious fact

Since the differences between rule-utilitarianism and rule-consequentialism are not relevant to the issues discussed here, I will refer to all of them as versions of rule-consequentialism.

[6] Ridge, 'Variable-Rate', p. 248.

[7] K. Tobia, 'Rule Consequentialism and the Problem of Partial Acceptance', *Ethical Theory and Moral Practice* 16 (2013), pp. 643–52, at 645.

[8] For example, Hooker and Fletcher identify a specific way in which this general problem can present itself when they point out that VRRC's calculation could be 'skewed by an anomaly'. B. Hooker and G. Fletcher, 'Variable *versus* Fixed Rate Rule Utilitarianism', *The Philosophical Quarterly* 58 (2008), pp. 344–52, at 349. See also, H. Smith, 'Measuring the Consequences of Rules', *Utilitas* 22 (2010), pp. 413–33, at 416.

[9] Hooker, *Ideal Code*, p. 84.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

that any halfway decent moral code would have very bad consequences if it were widely rejected. (Of course, just how bad the consequences of rejection would be depends on what is accepted in its place – if those who rejected a code did so in order to accept an even better code, then widespread rejection could actually increase a code's value.[10] To sidestep such complications, I shall assume that those who reject a code are immoralists.) By VRRC's criterion, the bad consequences of widespread rejection would count against a decent moral code just as much as the good consequences of widespread acceptance would count in its favour, but that is an unacceptable result. The fact that there would be dire consequences if a code were widely rejected certainly does not seem to be the sort of thing that should count *against* the code; if it should count for anything, it would seem to be a mark in the code's *favour*. To put this point in even sharper relief, consider a perverse moral code – say, one that requires continual acts of murder, revenge and abuse. In this case, the code's value will decrease as its acceptance rate increases, and its highest expected value will presumably be found where its acceptance rate approaches 0 per cent. But it is deeply counterintuitive to count the fact that we'd all be better off if no one accepted a code as a mark in the code's favour; it seems obvious that this fact should at least be given *less* weight than the perverse code's low value at higher acceptance rates, if it should be given any weight at all.

But perhaps the most obvious reason to deny that all acceptance rates should be granted equal weight is the fact that some acceptance rates are far more likely to obtain than others. If there were a reasonably high likelihood that a moral code could achieve 80 per cent acceptance and it is exceedingly unlikely that the code would achieve 100 per cent acceptance, then giving equal weight to the code's value at both rates would be unreasonable.[11]

So VRRC remains arbitrary in some respects, and more importantly, it fails to identify a relevant criterion for comparing the value of competing codes. Two more recent proposals have sought, in different ways, to rectify these difficulties.

Optimum Rate RC

Holly Smith follows Ridge in proposing that RC consider all acceptance levels, but instead of doing so in order to calculate a code's average expected value, she does so with an eye towards identifying the acceptance rate that *optimizes* each code's expected value. The expected value of a code at any given level of acceptance will be a function of both

[10] I am grateful to Max Parish for calling to my attention to the complexities involved in calculating the consequences of a code's rejection.

[11] Tobia, 'Partial Acceptance', p. 646.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

the benefits that accrue from that percentage of people accepting and generally complying with the rules and the costs involved in achieving that level of acceptance. Thus, for decent moral codes, expected value will increase as the acceptance rate increases, but at some point will begin to fall back off, as the additional costs of securing ever higher rates of acceptance begin to outweigh the added benefits. One code may have its highest expected value at 75 per cent acceptance, another code at 83 per cent, and so on. According to Optimum-Rate RC (ORRC), codes are to be compared to each other on the basis of their optimum acceptance rates.

> (ORRC) An *ideal code* is a code whose optimum acceptance rate is at least as high as the optimum acceptance rate of any alternative code.[12]

Unlike FRRC, ORRC does not select an acceptance rate arbitrarily, and unlike VRRC, it pays attention only to acceptance rates that intuitively seem to be relevant.

However, in spite of these advantages, Tobia contends that ORRC is not the best criterion for selecting codes.[13] ORRC seeks the code with the *potential* for the greatest benefit, but the code that *potentially* maximizes benefit could be a risky bet. That is, it could turn out that the code with the highest optimal rate would achieve such optimal outcomes only within a very narrow range of acceptance rates; moreover, ORRC does not take into consideration how likely it is that the code would achieve those optimal acceptance rates. The possibility that ORRC's ideal code might optimize consequences at an acceptance rate that is unlikely to actually obtain leads Tobia to propose an alternative criterion.

Maximizing Expectation Rate RC

Tobia's Maximizing Expectation Rate RC (MERRC) also looks at a code's value at every level of acceptance; however, instead of seeking to identify the code with the highest *potential* value, it seeks to identify the code with the highest likelihood of achieving a high value. In other words, MERRC seeks to identify the code with the highest *weighted average* value, each acceptance rate being weighted according to its likelihood of obtaining.

> (MERRC) An *ideal code* is a code whose weighted average value is at least as high as the weighted average value of any alternative code.[14]

[12] Smith, 'Measuring Consequences', p. 418. I have slightly modified Smith's wording for the sake of uniformity with the other formulations I am discussing.

[13] Tobia, 'Partial Acceptance', pp. 647–8.

[14] Tobia, 'Partial Acceptance', p. 649. My formulation is equivalent to Tobia's, but I have significantly modified his wording for the sake of clarity and uniformity with the

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

Like ORRC, MERRC avoids selecting a rate arbitrarily and offers an intuitively relevant criterion for comparing codes. But like the preceding theories, MERRC faces some difficulties.

MERRC's formulation requires that for each possible rate of acceptance (0–100) we should multiply the code's expected value at that rate by the likelihood of that rate's obtaining. However, these two values are both dependent upon a third factor in a way that renders this calculation problematic. There is no simple fact of the matter concerning how likely it is that a code will achieve a given acceptance rate, since the likelihood of acceptance depends on the methods and amount of effort put into teaching and enforcing the code. One method of teaching and enforcement may make an 80 per cent acceptance rate highly probable, while another method would make it improbable. So a code's likelihood of achieving any specific acceptance rate will vary depending upon the methods of teaching and enforcement. The code's expected value will also vary depending on this factor. After all, the method and amount of effort required to teach and enforce a code are among the costs that must be subtracted from the expected benefits in order to calculate a code's overall expected value. So both elements in MERRC's calculations are indeterminate, and both vary depending on the methods and effort put into teaching and enforcement.

Furthermore, there is no obvious way to stipulate a solution to this difficulty. We cannot simply require that the calculations be made assuming the method of teaching and enforcement most likely to achieve high levels of acceptance, since that method may be prohibitively costly. Nor can we require that the calculations be made assuming the least costly method of teaching and enforcement, since this method may not be likely to achieve very high levels of acceptance. What we presumably would want is the method that would achieve the best balance of low cost and high acceptance; but there is no way to say in advance what method that would be, nor can we be confident that it would be the same method for each possible acceptance level. In spite of the intuitive appeal of its criterion, it is far from clear that the calculations MERRC requires could actually be made to work.

Uniform vs. non-uniform acceptance

We have seen that there are unique difficulties facing each formulation of RC we have considered so far. However, there is a further difficulty that applies to all four formulations. Whether they arbitrarily stipulate

other formulations I am discussing. Tobia formulates his criterion as a weighted *sum*, rather than a weighted *average*; restating MERRC in terms of a weighted *average* helps to clarify the respects in which it is both similar and dissimilar to VRRC.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

an acceptance rate, seek the optimum acceptance rate, or compute a simple or weighted average value for all acceptance rates, all of the formulations assume that it makes sense to evaluate the values of codes based upon the assumption of a uniform rate of acceptance for the entire code. That is, they all assume that it makes sense to ask questions of the form: *what would be the expected value of code A on the assumption of N per cent acceptance?* In reality, the various rules included within a code will almost always achieve unequal acceptance rates. This is not just an unfortunate complication that can safely be glossed over for theoretical purposes, because moral codes will almost certainly have lower expected values on the assumption of a uniform acceptance rate than they would achieve on the assumption of varying acceptance rates for different rules.

Compare, for example, the rules against murder and illegally downloading copyrighted music. According to a quick Google search, roughly 60 per cent of Americans illegally download music, while roughly .0048 per cent commit murder. These are, of course, rates of compliance rather than acceptance. It may be that some of those who illegally download music fail to comply with a rule they accept, and some percentage of the population may comply with the rule against murder purely out of fear of the consequences of being caught and not because they accept the rule. Nevertheless, such considerations would do little to close the broad gap between these rates, and so for present purposes I will treat them as if they were acceptance rates. If a mere 40 per cent of Americans currently accept the rule against illegal downloading, it would probably take a tremendous amount of effort to achieve a 90 per cent acceptance rate for that rule – so much effort that the necessary teaching and enforcement costs would probably outweigh the benefits. For that rule a 90 per cent acceptance rate would probably be impractical, even if something higher than 40 per cent would be more ideal. By contrast, if the acceptance rate for the rule against murder were a mere 90 per cent, the results would be horrific. Any code that included these two rules would have dramatically worse consequences on the assumption of a uniform 90 per cent acceptance rate than it would on the assumption of non-uniform rates of, say, 70 per cent (copyright laws) and 99.998 per cent (murder).

It might be possible, at least in principle, to extend some of the theories we have considered to account for non-uniform acceptance rates. For example, a variant of ORRC could calculate the value of a code for each possible combination of non-uniform acceptance rates in order to identify the combination of acceptance rates that optimizes the code's value. However, while such extensions might work in theory, they would not be attractive options. We should want a formulation of

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

RC to offer not only a formal characterization of the ideal code, but also a practically useful way of thinking about competing codes in order to determine which satisfies the formal characterization. Even if a non-uniform version of one of these theories could offer a satisfactory formal characterization of the ideal code, the calculations involved would be so numerous and exceedingly complex that it would be of no practical value in identifying the ideal code.[15] The point is not merely that the calculations are too complex to be carried out with any degree of precision. That is bound to be true of all formulations of RC. Rather, the point is that the calculations involved are so complicated that we would be hard-pressed even to make educated guesses concerning the values of a code at various sets of non-uniform acceptance rates, and so the extended theories would not offer any practical guidance at all.

A NEW FORMULATION: CALCULATED RATES RC

Stipulated vs. calculated rates

Reflecting on the ways we naturally think about the consequences of rules in everyday life suggests an alternative approach to formulating RC that has considerable advantages over the theories considered above. Let me begin with a simple pedagogical example. It is a sort of rule in all of my courses that students should come to class prepared, which typically means they should have read the assigned readings carefully and thoughtfully. Of course, I know that only a handful of my most dedicated students will consistently follow this rule unless my syllabi specify methods of encouraging and enforcing it. My course requirements typically include assignments intended to encourage and develop good reading habits, and the fact that these assignments factor into students' course grades provides a sort of built-in enforcement mechanism. I have experimented with various types of assignments: pop reading quizzes, online discussion forums, short essays on issues related to readings, and so on. My goals in choosing assignments are analogous to RC's goals in selecting a moral code: I want to maximize the benefits of student acceptance of and compliance with the rule about coming to class prepared, but I want to do so in a way that minimizes the costs (both to myself and my students) involved in achieving those benefits. I consider various methods of encouragement and enforcement (in this case, various types of assignments and grade incentives), and then very roughly estimate their expected values. But in order to do that, I have to make estimates about the levels of

[15] Hooker and Fletcher raise a similar objection to Ridge's VRRC. See Hooker and Fletcher, 'Variable *versus* Fixed', pp. 349–50.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

acceptance and compliance I can expect to achieve with each method. If an assignment is too demanding, I anticipate that more students will give up in despair. An easier assignment may achieve higher compliance, although it may not achieve the same level of careful reading. I also consider the costs involved with each method. How much additional time will it take for students to complete? How much extra grading work will it create for me? My goal is to select the method with the best value – the best student performance for the least extra work.

Notice how this thought process differs from the approach involved in the Fixed, Variable, Optimum and Maximal Expectation Rate formulations of RC. I do not approach my pedagogical problem by stipulating some target rate of acceptance and then calculating the benefits minus the costs involved in achieving that rate. Instead, I stipulate some method of teaching and enforcement and then calculate the expected costs and benefits, including the rates of acceptance and compliance, I can expect those methods to achieve.

This example highlights the fact that there are at least three interrelated variables in rule-consequentialist calculations: (a) acceptance rate, (b) compliance rate, and (c) methods of teaching and enforcement. A value for one or more of these variables has to be stipulated as an input to the calculation in order to estimate the value of the remaining variables. The Fixed, Variable, Optimum and Maximizing Expectation Rate formulations all calculate a code's value based upon *stipulated* rates of acceptance. The stipulated acceptance rate is then used to calculate roughly the code's expected costs and benefits, including the costs of teaching and enforcing the code in a way that would achieve that level of acceptance.

In spite of its prevalence in the RC literature, the stipulated rate approach is not at all natural or intuitive. We certainly do not approach familiar debates about moral issues in public policy this way. Does the death penalty deter serious crime better than alternative punishments such as life in prison? Would we better off if we legalized marijuana use? What is the best way to deal with copyright infringement? Should we allow same-sex couples to marry? Should we legalize physician-assisted suicide, and if so, how could it best be regulated? In such debates, rational disputants often find themselves comparing competing rules and competing methods of teaching and enforcing those rules, looking for the rules and methods that would achieve the best results. Our common-sense approach to such problems suggests the possibility of alternative formulations of RC in which the methods of teaching and enforcement are stipulated, and the acceptance and compliance rates are calculated as part of the consequences of implementing those methods.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

A sketch of Calculated Rates RC

A Calculated Rates formulation of RC evaluates not only codes or sets of rules, but also stipulated methods of teaching and enforcement. To avoid any confusion, we can continue to use the term *code* to refer simply to a complete set of moral rules, and we can refer to a *code + stipulated methods of teaching and enforcement* as an *expanded code*. Calculated Rates RC (CRRC) can then be formulated as follows:

> (CRRC) An *ideal code* is a *code* included in an *expanded code* whose value is at least as high as the value of any alternative *expanded code*.

In calculating the expected value of an expanded code, we would have to make realistic estimates about the levels of acceptance and compliance we could expect its teaching and enforcement methods to achieve. I have already suggested that this would bring RC more in line with our common-sense approach to moral problems, and I regard this as a significant advantage in itself. I now want briefly to consider two additional advantages of CRRC over stipulated rate formulations.

First, unlike stipulated rate formulations, CRRC does not unrealistically assume uniform rates of either acceptance or compliance. If we begin with a moral code and then stipulate a proposed method of teaching and enforcing it, we will expect that various elements of the code will achieve differing rates of acceptance. We can predict that a rule forbidding murder will more easily achieve a high rate of acceptance and compliance than a rule demanding some form of altruistic behaviour such as giving to charitable causes. This is a major advantage of CRRC because achieving high compliance rates is far more important for some rules than it is for others. Generally speaking, the more severe the costs of non-compliance, the more it matters that we achieve high rates of compliance and acceptance. We are willing to go to greater lengths to ensure compliance with the rule forbidding murder than we are to ensure compliance with, say, copyright laws. At the other end of things, there may be rules worth having, even though the cost of non-compliance with them is so low that strict enforcement methods would be unjustified. It is certainly worthwhile to try to gain widespread acceptance of the rule that we all ought to recycle as much as reasonably possible; and yet, the costs of a single violation of this rule are typically so low that strict enforcement would be unjustified. In short, the rules included in an ideal code will be taught and enforced with varying degrees of stringency. This is an important consideration that CRRC naturally factors into its evaluation of codes, while uniform rate formulations ignore it.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

Second, whereas stipulated rate formulations must decide whether to stipulate a rate of compliance or of acceptance, CRRC is not forced to make such a choice. These rates are obviously related to each other. Increased internalized acceptance of a rule will almost always increase compliance with it. But their relationship is also complex. Sometimes we comply with rules we do not accept; for example, we may obey speed limits even though we regard them as unnecessarily restrictive if we know that that they are strictly enforced and do not want to deal with the consequences of being ticketed. While it is preferable to get people to comply with rules because they accept them, it is often crucial to have methods in place to motivate compliance among those who do not accept a rule, so part of an expanded code's value will depend on its ability to motivate compliance among the non-accepting. Similarly, we sometimes fail to comply with rules even though we do accept them. Since a high acceptance rate is no guarantee of a comparably high compliance rate, methods of motivating compliance may be crucial even for rules that we expect to achieve high acceptance rates.

Stipulated rate formulations are at a disadvantage when it comes to such complexities. *Compliance* formulations are especially hampered, because a stipulated compliance rate tells us nothing about how that level of compliance is achieved. Do people comply because they have accepted the rule, or only because they are coerced by threats of punishments for non-compliance? Stipulating a compliance rate will determine rather little about either of the other two variables. Stipulated acceptance rate formulations are somewhat better off, since in many cases acceptance is a primary factor determining compliance. To the extent that compliance is a function of acceptance, a stipulated acceptance rate will also give a rough sense of what rate of compliance can be expected. However, in cases where there are strong incentives for individuals to violate rules they accept, the acceptance rate will still underdetermine the compliance rate. In such cases, predicting the likely rate of compliance requires information about how the rules are to be enforced in order to determine whether those methods would provide adequate extra motivation. So although stipulated rate formulations can, to some extent, justify methods of dealing with these sorts of complexities,[16] it is unlikely that they can do so as effectively as a formulation that makes realistic predictions about both rates. Since CRRC stipulates a code and a method of teaching and enforcing the code, it can make realistic predictions about both acceptance and compliance and is not forced to focus on one at the potential expense of the other.

[16] Much of Hooker's case for formulating RC in terms of less than complete acceptance is based in consideration of these sorts of complexities; see Hooker, *Ideal Code*, pp. 80–5.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

The most obvious weakness of CRRC is that it will often be difficult to predict what level of acceptance and compliance an expanded code would achieve; as a result, it will often be difficult to calculate precisely an expanded code's expected value. Although this is undeniable, it is not a compelling reason to reject CRRC in favour of a stipulated rate formulation because it is equally difficult to predict just how much cost would be involved in achieving a stipulated rate of acceptance. No matter what formulation is adopted, RC has to make the best of working with rough and uncertain estimates.

Of course, all this response shows is that CRRC is no worse off in this respect than any other formulation of RC. One might think that this is a poor consolation if CRRC is also no better off than other formulations of RC; that is, one might reasonably worry that the difficulty of such calculations calls into question the viability of all formulations of RC. I think rule-consequentialists should take this concern seriously. Earlier, I contended that an adequate formulation of RC should not only offer a formal characterization of the ideal code, but should also offer practical guidance concerning which code satisfies that characterization. In so far as formulations of RC focus on the evaluation of entire codes, it is difficult to see how they can be of much help in offering the sort of practical guidance we want. This problem is especially pressing when we try to imagine applying the theory to codes that differ significantly from conventional morality – the more distant the code, the less confident we can be in our predictions of its expected consequences.

Richard Brandt has offered what I consider the most promising approach to overcoming this difficulty:

> [O]ne can argue that the *existing* moral codes are time-tested: there are socio-psychological forces which tend to prune away pointless prohibitions, and to some extent to bring about requirements suited to new problems (e.g. about conservation of resources). So one can say that what the intelligent rule utilitarian wants, in seeking to identify the optimal code, is primarily to expedite this process.[17]

In other words, instead of evaluating entire codes, in practice RC should focus on particular rules (or related subsets of rules). Another way to put this point is to say that RC can focus on comparing the code that represents conventional morality with alternative codes that differ from conventional morality in some specific respect. Brandt describes his approach as 'incrementalist' because rather than seeking to identify the ideal code all at once (from among an infinite number

[17] R. Brandt, 'Fairness to Indirect Optimific Theories in Ethics', *Ethics* 98 (1988), pp. 341–60, at 356–7.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press

of possible options), it works towards an ideal code one small step at a time. Brandt's incrementalism makes the calculations involved in CRRC much more manageable. After all, we already have (or can acquire) empirical evidence regarding the non-uniform acceptance and compliance rates of the rules of conventional morality. So moral debate will typically end up focusing on what we could expect to happen if conventional morality were supplemented with some additional rule or if some part of conventional morality were replaced with an alternative rule. In addition to making the calculations involved much more manageable, incrementalism brings RC more in line with the actual practice of moral debate and the reality that moral progress is usually accomplished through a slow and gradual process.

CONCLUSION

In this article I have called attention to two previously unnoticed distinctions within possible formulations of RC: the distinction between *uniform* and *non-uniform* rate formulations and the distinction between *stipulated* and *calculated* rate formulations. Previous formulations of RC have tended to be *uniform stipulated* rate formulations. The new formulation I have proposed, CRRC, is a *non-uniform calculated* rate formulation. Instead of stipulating uniform acceptance rates as inputs to our calculations, CRRC stipulates methods of teaching and enforcement and leaves the non-uniform acceptance rates among the variables we solve for in our calculations. With respect to the problem of making realistic assumptions about partial acceptance and compliance rates, CRRC has considerable advantages over stipulated rate formulations and, as far as I can tell, no disadvantages that those theories do not share. As it turns out, the solution to the acceptance rate problem is simply to solve for the acceptance rates.[18]

tmiller@leeuniversity.edu

[18] I am grateful to Christopher Stephenson and my audience at the University of Oklahoma for comments on earlier versions of this material.

https://doi.org/10.1017/S0953820815000229 Published online by Cambridge University Press