

Michigan Law Review

Volume 80 | Issue 4

1982

Utilitarianism Reformed

L. W. Sumner
University of Toronto

Follow this and additional works at: <https://repository.law.umich.edu/mlr>



Part of the [Law and Philosophy Commons](#)

Recommended Citation

L. W. Sumner, *Utilitarianism Reformed*, 80 MICH. L. REV. 701 (1982).

Available at: <https://repository.law.umich.edu/mlr/vol80/iss4/20>

This Review is brought to you for free and open access by the Michigan Law Review at University of Michigan Law School Scholarship Repository. It has been accepted for inclusion in Michigan Law Review by an authorized editor of University of Michigan Law School Scholarship Repository. For more information, please contact mlaw.repository@umich.edu.

UTILITARIANISM REFORMED

*L. W. Sumner**

UTILITARIANISM AND CO-OPERATION. By *Donald H. Regan*. New York: Oxford University Press. 1980. Pp. xiv, 279. Cloth \$37.50; paper \$15.95.

In 1780, the first edition was printed of Jeremy Bentham's *Introduction to the Principles of Morals and Legislation*, the earliest modern statement of the moral theory that has come to be known as utilitarianism.¹ Despite the vicissitudes of philosophical fashion since Bentham's time, the theory has seldom relinquished its position as the centerpiece of Anglo-American moral, political, and legal philosophy. The issues that have animated the enormous literature on utilitarianism may be broadly grouped into two categories: the external and the internal. The former concerns the relations between utilitarianism and rival theories, or between utilitarianism and "common sense." The critical question here is the adequacy of the theory as a general guide to individual moral decision, social choice, and the construction of political and legal institutions. This external debate has been paralleled by an equally vigorous intratheoretical inquiry carried on within the utilitarian camp itself. The critical issue here is the comparative assessment of the various versions of the theory. Utilitarianism, like most political and religious ideologies, appears to be a single unified doctrine only from afar. It is, in fact, a family of theories loosely linked by a shared set of basic presuppositions. The internal debate concerning the relative merits of the many possible members of this family has not been entirely insulated from the external one, for some versions of the theory have been devised and defended for their supposed ability to surmount what have been felt to be formidable substantive difficulties. But much of the discussion has assumed the adequacy of a generally util-

* Professor of Philosophy, University of Toronto. B.A. 1962, University of Toronto; M.A. 1964, Ph.D. 1965, Princeton University. The author has recently published a book on *Abortion and Moral Theory* (1981). — Ed.

1. Actual publication was delayed until 1789 owing, Bentham tells us, to his discovery of some unexpected metaphysical difficulties. His attempt to resolve them led to the composition of *Of Laws in General*. See J. BENTHAM, *OF LAWS IN GENERAL* (H.L.A. Hart ed. 1970) (the manuscript was first published in 1945 under the title *The Limits of Jurisprudence Defined*).

itarian approach and has been directed toward identifying the variant that best meets the theory's own standards.

The internal issues then subdivide into two categories. The first concerns the content of the theory: What determines the value of an action's consequences? The second concerns the theory's form: What is the connection between the moral properties of an action and the value of its consequences? Even a limited number of plausible options concerning each of these elements will combine to generate a sizable variety of theories, each of which may reasonably be represented as utilitarian.

Donald H. Regan's *Utilitarianism and Co-operation*, published exactly two hundred years after the printing of Bentham's pioneering work, provides fresh evidence that the theory's resources have not yet been exhausted. The book falls nicely into the second subdivision of the internal debate. Regan does not try to defend utilitarianism against rival approaches; nor does he tell us what to count as utility. He attempts instead to assess the relative merits, from a utilitarian point of view, of some of the theory's traditional forms: act-utilitarianism, rule-utilitarianism, and utilitarian generalization.

Regan's arguments and conclusions are important, and I will return to them shortly. First, however, I have some cautions for potential readers. Despite Regan's academic post at the University of Michigan Law School, the book contains no legal theory whatsoever. Regan confines himself to utilitarianism as a moral theory and ignores its political and legal implications. The discussion is, moreover, highly theoretical. Regan is concerned with practical moral problems only to the extent necessary to expose the strengths and weaknesses of possible versions of the theory. Most of the cases that he uses to test the theory's traditional varieties are themselves highly formal. Although he pays some attention to real-life examples illustrating the various theories, he never carries the analysis past the point where the desired theoretical conclusions have emerged. This avoidance of both political-legal and practical-moral issues means that the book will interest only moral philosophers.

Beyond this initial narrowing of the potential audience lie some further barriers. The questions that Regan addresses are extremely intricate, and an adequate treatment of them cannot fail to be itself intricate. Even the professional philosopher will find the book dense and difficult. This is not necessarily a defect. Indeed, Regan has done his best to mitigate the reader's inevitable labors in persevering through the argument. The main track of his argument is not hard to follow if one passes by the side trails that are now and again ex-

plored. The book begins with a map of the route and signposts are located at frequent intervals. Complications and digressions are clearly labeled and often segregated from the central argument. Finally, not the least of the book's virtues is that the author has confined technical vocabulary to the minimum necessary for efficient deployment of the argument, and he writes in plain and accessible English. In short, as many concessions have been made to the reader as are compatible with analytic rigor. But the issues remain complex, and their exploration forbidding.

Moreover, many moral philosophers, who are, after all, trained to read hard books in their subject, will find the issues unfamiliar. In the past two decades, the literature on the varieties of utilitarianism has undergone a dramatic population explosion. The debate that Regan has joined has an extensive history, and an understanding of his contribution is aided considerably by familiarity with that history. Even though Regan has done much to explain the issues to the uninitiated, a reader's lack of background will be an added liability that he will need to surmount.

Before the book's potential audience shrinks to the vanishing point, let me provide some incentives for tackling this book. The literature on this subject, as on most, consists of a very large number of minor pieces and a very small number of landmark works. The two principal landmarks to date have been David Lyons's *Forms and Limits of Utilitarianism*² and D.H. Hodgson's *Consequences of Utilitarianism*.³ Now there is a third. (It is an interesting sidelight that all three books began life as doctoral dissertations; apparently this is the sort of work of which one is capable only during one's graduate studies.) Regan's book has raised the state of the art to a new level. It is rigorous, lucid, comprehensive, and original. It illuminates virtually every issue it addresses, including those that it shares with Lyons and Hodgson. A later landmark, of course, profits from the opportunity to criticize its predecessors, but even adjusting for this historical advantage, Regan's work is an impressive accomplishment. No one interested in the structure of utilitarian theories can afford to ignore it.

* * *

In the remainder of this Review I will first briefly outline Regan's main argument and then join some of the issues that it raises.⁴ His

2. D. LYONS, *FORMS AND LIMITS OF UTILITARIANISM* (1965).

3. D.H. HODGSON, *CONSEQUENCES OF UTILITARIANISM* (1967).

4. In my formulation of the argument, I have made minor alterations in some of Regan's definitions and theories, in the interest both of greater uniformity among them and of economy

objective is to compare the merits of three forms of utilitarianism. The most familiar of these is act-utilitarianism:

AU An act is right if and only if it produces the best possible consequences in the situation in which it is performed.⁵

Of the various possible versions of rule-utilitarianism and utilitarian generalization, Regan selects (for reasons that will shortly become clear) what he calls the "co-ordinated optimization principle":

COP An act is right if and only if it is prescribed (for the agent whose act is in question) by that universal prescription for action, the universal satisfaction of which would produce the best possible consequences.

The third member of the trio is what Regan calls "co-operative utilitarianism":

CU Each agent ought to co-operate, with whoever else is co-operating, in the production of the best possible consequences, given the behavior of nonco-operators.

The first step of the argument is a definition of what it is for a moral theory to be satisfied:

Theory T is satisfied when an agent actually does what T requires him to do in a given situation.

It is then possible to define two properties that a particular moral theory might possess:

Theory T has PropAU if any agent satisfying T in any situation produces by his act the best possible consequences.

Theory T has PropCOP if all agents satisfying T in all situations jointly produce by their acts the best possible consequences.

Regan suggests that the fundamental intuitive idea underlying act-utilitarianism is that a moral theory ought to be a good theory for *each individual* to follow (*i.e.*, it ought to have PropAU), and that the fundamental intuitive idea underlying the traditional forms of rule-utilitarianism and utilitarian generalization is that a moral theory

of vocabulary. Any resulting distortions of his meaning are, of course, entirely my responsibility.

5. In AU, and in all subsequent contexts, "best possible consequences" must be read to mean "consequences at least as good as would be produced by any alternative." See Regan's formulation of AU, p. 12.

ought to be a good theory for *all individuals* to follow (*i.e.*, it ought to have PropCOP). Ideally, then, a utilitarian should want a moral theory to have both properties. We thus have two standards, internal to utilitarianism, for judging versions of the theory.

Regan’s main substantive contentions may now be readily enumerated:

- (1) AU has PropAU but not PropCOP.
- (2) COP has PropCOP but not PropAU.
- (3) No form of rule-utilitarianism or utilitarian generalization has PropAU (and most do not even have PropCOP).
- (4) No form of utilitarianism that resembles AU and COP in being “exclusively act-oriented” can have both PropAU and PropCOP.
- (5) CU, which is not exclusively act-oriented, has both PropAU and PropCOP.
- (6) Therefore, on internal utilitarian grounds, CU is preferable to AU and also to any form of rule-utilitarianism or utilitarian generalization.

For these claims to be fully intelligible we need a definition of “exclusively act-oriented.” I will return to this later.

Although the main line of Regan’s argument can be quickly sketched, a detailed description is a much more complicated matter. I cannot hope to do justice to more than a small part of it. I shall therefore pass over Regan’s critique of rule-utilitarianism and utilitarian generalization, and focus on the relative merits of AU and CU.

That AU has PropAU is obvious. That it does not have PropCOP is shown by the following case. Imagine two persons, Row and Column, each of whom is presented with two possible options: for Row these are r1 and r2 and for Column they are c1 and c2. They must choose between these options independently; neither is able to communicate with, or in any way to influence the decision process of, the other. Their actual choices will jointly produce an outcome, the values of the possible outcomes being given in the following matrix:

		Column	
		c1	c2
Case 1	Row	r1	10 0
		r2	0 6

Let us say that an outcome is *optimal* when it is best (*i.e.*, at least as good as any other outcome). In this situation, r1c1 is the optimal outcome. If Row and Column by their choices jointly produce that outcome, then each satisfies AU: given Row's r1, AU requires Column's c1, and vice versa. But if Row and Column jointly produce r2c2 then each still satisfies AU: given Row's r2, AU requires Column's c2, and vice versa. AU is thus universally satisfied either by r1c1 (the optimal outcome) or by r2c2 (a suboptimal outcome); AU is *indeterminate* between these outcomes. Because universal satisfaction of AU does not ensure joint production of the best consequences, AU does not have PropCOP.

In Case 1, no stipulation was made of either Row's or Column's reasons for choice, but we are free to imagine that they are both act-utilitarians. Let us say that an outcome is an *equilibrium* when each agent's choice produces the best possible consequences, given the actual choice of the other; let us also say that the agents have *coordinated* when they have jointly produced an equilibrium outcome. Act-utilitarian agents will wish to coordinate, since they will each have failed to satisfy AU if they jointly produce a nonequilibrium outcome. But universal satisfaction of AU requires only that they coordinate, not that they coordinate to produce the optimal outcome.

Case 1 establishes that AU lacks PropCOP. Since this is the defect in AU that leads Regan to prefer CU, it is important to ascertain just how serious the defect is. Regan has not shown that AU is self-defeating. (A moral theory is self-defeating when its satisfaction impairs the achievement of its own aims.)⁶ AU is not self-defeating for two reasons. First, it does not fail according to its own standards. PropAU captures the fundamental intuitive idea behind AU, and AU has PropAU. AU fails only by the standards of rule-utilitarianism and utilitarian generalization. These may, of course, be standards worth meeting, and they are utilitarian standards, but they are not act-utilitarian standards. Second, and more important, while it is true that universal satisfaction of AU *may* produce a suboptimal outcome, it is not true that it *must* produce such an outcome.⁷ The

6. For a useful discussion of self-defeatingness in moral theories, see Parfit, *Is Common-Sense Morality Self-Defeating?*, 76 J. PHIL. 533 (1979).

7. See pp. 54-65. Here lies the crucial difference between AU and some other consequentialist principles. AU directs an agent to produce by his act the best *overall* consequences for everyone; it thus assigns each agent a common goal. What we might call act-egoism (AE) directs an agent to produce by his act the best *personal* consequences for himself; it thus assigns each agent a different goal. AU may face coordination problems in which universal satisfaction of the theory is compatible with a suboptimal outcome. But AE, besides facing analogous problems (in situations of perfect identity of interest), may also face Prisoners' Dilemmas (in situations of partial conflict of interest) in which universal satisfaction of the theory ensures a

optimal outcome in act-utilitarian coordination problems, after all, is also always an equilibrium. Production by each agent individually of best consequences is therefore a necessary, though not a sufficient, condition of production by all agents collectively of best consequences.

Even if AU is not self-defeating, however, the fact that a set of act-utilitarians may each satisfy their theory while coordinating on a suboptimal outcome still seems bothersome. The degree of bother varies in direct proportion to the size of the gap between the optimal and suboptimal outcomes. Consider the following matrix, for all values of $n > 1$:

			Column	
			c1	c2
Case 2	Row	r1	n	0
		r2	0	$1/n$

Or again:

			Column	
			c1	c2
Case 3	Row	r1	10^n	0
		r2	0	10^{-n}

The second, and especially the third case show that as the value of n increases it becomes correspondingly more important that act-utilitarians coordinate to produce the optimal outcome.

Regan's argument does not show that they always, or indeed ever, fail to coordinate. This is its most important limitation. For all that the argument tells us, sets of act-utilitarian agents might invariably succeed, as a matter of fact, in coordinating to produce the optimal outcome. The indeterminacy of AU in Cases 1-3 may easily mislead us: AU's indeterminacy *only* means that it is universally satisfied by both the optimal and the suboptimal outcome; it does *not* mean that AU requires agents to be indifferent between the two outcomes. If Row and Column in Cases 1-3 choose $r1c1$, they jointly produce better consequences than if they had chosen $r2c2$, but each also individually produces better consequences than if they had cho-

suboptimal outcome. There is no counterpart for AU of conflict-of-interest choice situations. See Parfit, *supra* note 6.

sen r2c2.⁸ Since AU enjoins them to produce the best consequences, each, therefore, has an AU-given reason for preferring that they coordinate to produce the optimum. Neither, of course, can bring about this result unilaterally, but each has a reason for doing what he can to help bring it about. This fact explains why we should find it inexplicable, indeed perverse, if the two act-utilitarians in Cases 1-3 were to coordinate on the suboptimal outcome.⁹ If coordinating to produce the optimal outcome were just as easy as producing the suboptimal outcome, what reason could either have for unilaterally defeating that result?

In coordination problems, one equilibrium result may possess some property that makes it stand out from the others; it may, that is, be *salient*.¹⁰ It then becomes an obvious point on which agents who wish to coordinate may converge. In the first three Cases, r1c1 stands out simply because it is optimal. This fact provides act-utilitarians with an AU-given reason for trying to coordinate to produce that outcome. They have available to them a strategy which, if universally followed, will ensure such coordination. Consider Case 1 again. Row's r1 will yield (depending on Column's choice) either 10 or 0; call this his 10-0 option. His r2 will yield either 0 or 6; call this his 0-6 option. Column also faces 10-0 and 0-6 options. Row and Column are thus symmetrically situated.¹¹ Neither option is (even weakly) dominant; however, they do have different maxima. If Row and Column both choose the option with the higher maximum (if both "maximax") they will coordinate to produce the optimum. And they will do so in all act-utilitarian coordination problems where there is a unique optimal result. Nothing, incidentally, depends on the fact that Row and Column are symmetrically situated in the first three Cases. This symmetry is absent in the following case:

		Column	
		c1	c2
Case 4	Row	r1	10 5
		r2	0 6

8. I use here the marginal conception of consequences (correctly) favored by Regan, pp. 13-17.

9. It is also what lies behind the obviousness argument that Regan dismisses. See pp. 21-23.

10. I have borrowed this term from Gauthier, *Coordination*, 14 DIALOGUE 195 (1975). The classic discussion is T. SCHELLING, *THE STRATEGY OF CONFLICT* (1960).

11. By saying that Row and Column are symmetrically situated I mean that the possible (individual) consequences of their options are identical. Regan uses a different notion of symmetry. See pp. 23-25.

Here Row faces 10-5 and 0-6 options while Column faces 10-0 and 5-6 options. Nonetheless, if both maximax they will coordinate to produce the optimal result, r1c1.

Maximax is, therefore, a rule or procedure for choosing among options that will naturally suggest itself for universal adoption by act-utilitarians in coordination problems. What is its status relative to AU itself? We should remember that satisfying AU generally requires considerable information about the choice situation. AU is difficult to satisfy when that information is, for whatever reason, incomplete. Consider cases of "individual choice under uncertainty." Here, the outcome is determined partly by the individual's choice and partly, not by the independent choice of another agent, but by nature. Further, the agent lacks sufficient information even to assign objective probabilities to the various possible states of nature. Imagine that Row is told that an urn contains 100 balls, some red and some black, but is not told the ratio of one to the other. One ball is to be drawn at random from the urn and Row is required to choose between betting on red or black, with the following payoffs:

		Nature	
		Red	Black
Case 5	Row	Red	10 0
		Black	0 6

Assuming that Row wishes to produce the best consequences, how is he to choose? Given that he is confronted by 10-0 and 0-6 options, it seems reasonable for him to maximax, and thus to bet red. Maximax is a strategy or procedure intended to increase the likelihood of satisfying AU in a situation whose incomplete description makes it impossible to satisfy AU directly. Whether the strategy is best in a particular case will be determined by the way things actually turn out (whether AU is satisfied). Whether it is best in all cases of this sort will be determined by the way things actually turn out in the long run (how often AU is satisfied). AU thus remains the final test of the value of the strategy. In Case 5, maximax seems to be a sensible act-utilitarian strategy. This is, however, not always so in situations of "individual choice under uncertainty." Consider the following case:

		Nature		
		Red	Black	
Case 6	Row	Red	10	0
	Black	8	9	

Here maximax still dictates betting red, but betting black seems the wiser choice. Agents in situations of "individual choice under uncertainty" who are attempting to satisfy AU will have to attend to both the maxima and the minima of their options.

Whereas maximax is a fallible strategy for an individual act-utilitarian in choice under uncertainty, it is an infallible strategy for a group of act-utilitarians in coordination problems with a unique optimum. This very fact renders it salient for the members of such a group and thus increases the likelihood that they will all follow it and so will all coordinate to produce the optimal result.

An outcome may be salient for reasons other than its being an optimum. The bare schemata of coordination problems largely suppress the further properties that can make outcomes salient. In real-life problems, however, some of these additional cues are likely to be available to agents who are trying to coordinate. To the extent that they are available in a given case, act-utilitarians again increase their chances of successful coordination.

These further salient properties become particularly important when we discard a feature that has been shared by all cases thus far considered: their possession of a unique optimum.

		Column		
		c1	c2	
Case 7	Row	r1	10	0
	r2	0	10	

With no further cues, Row and Column have no reason for preferring r1c1 over r2c2; maximax does not determine a unique choice. Row and Column are therefore in danger of failing to coordinate on either equilibrium, and thus of failing to satisfy AU. Each has an AU-given reason for trying to coordinate, but no reason for aiming at a particular equilibrium outcome. If they are presented only with the abstract schema and are still prevented from communicating, they may have to resort to a randomized strategy that depends for its success upon luck. In cases with multiple optima, AU remains indeterminate in outcome since it is universally satisfied by any equilib-

rium, including a suboptimal one. But it is also indeterminate in strategy: because there is no unique optimum, there is also no salient choice procedure on which to converge. In real-life coordination problems with multiple optima, act-utilitarians who are unable to communicate will need to rely on further cues to render one of the optima salient.¹²

If, unlike Row and Column, they are able to communicate, they will have additional resources for ensuring coordination. In the first four Cases they should agree to maximax, and in Case 7 they should agree to choose one of the two equally attractive optima. Thus again, once the artificial constraints of the abstract problem are discarded, the chances improve that act-utilitarians will in fact manage to produce an optimal outcome.¹³

In some real-life coordination problems, communication and agreement are admittedly either impossible or excessively costly — whether because of the number of agents involved, or their spatial or temporal distance from one another, or whatever. In such cases, a further device is available to produce coordination: the establishment and enforcement of a rule. Navigation is an obvious example. Perhaps in each separate instance ships approaching one another would manage to agree to pass either to port or to starboard, but it is more efficient to select one or the other (perhaps arbitrarily), enforce it as a general rule, and punish violators. Society exists, in part at least, to solve coordination problems. Sanctions alter the consequences of the choices confronting the agents. In Case 7, for instance, enforcing a rule requiring Row's r1 and Column's c1 — agents choosing other options would be punished — would reduce the value of r2c2, thus converting Case 7 into something more like Case 1. And, of course, the fact that a given action is legally (or socially) required may itself render the action salient, even apart from its effect on the value of its consequences.

12. Regan's cooperative utilitarians must also resort to such cues in choice situations of this sort. *See* pp. 190-206.

13. Regan's discussion of communication and agreement among act-utilitarian agents, pp. 32-43, is the least convincing in the book. The issue is not how act-utilitarians develop practices such as language use or agreement keeping, but whether act-utilitarians who have such practices can use them to coordinate. The very existence of the practices renders certain procedures (such as agreeing to maximax and then keeping the agreement) salient; therefore, they increase the likelihood of optimal coordination (which is what act-utilitarians will be trying to achieve). Regan recognizes all this when he summarizes his objection: "[A]ll I claim is that it does not follow from the fact that certain parties are act-utilitarians that an 'agreement' between them will have any effect on their expectations or behaviour. Whether an agreement is effective or not depends on facts about the parties other than the fact that they are act-utilitarians." P. 37. These further facts include their ability to speak a common language and to make and keep agreements, properties generally possessed by real-life act-utilitarians.

Because AU does not have PropCOP, universal satisfaction of AU does not ensure coordination on an optimum. But universal satisfaction of AU is always compatible with coordination on some optimum, and act-utilitarians have an AU-given reason for wishing so to coordinate. Act-utilitarians, moreover, have a variety of procedures, strategies, and devices — salience, communication, agreement, enforcement — available for increasing the likelihood of effective coordination. AU's lack of PropCOP may, therefore, be a theoretical defect with few practical costs. It will not much matter in practice that universal satisfaction of AU does not *ensure* an optimal outcome if universal satisfaction of AU *in fact* generally results in an optimal outcome.

Still, the fact remains that universal satisfaction of Regan's CU does ensure production of an optimum. Is not a theory whose correct application by all guarantees an optimum superior, in practical terms, to one whose correct application by all might produce a suboptimum? Not necessarily. We need to distinguish two sorts of consequences that may flow from a theory's universal satisfaction. The first are the consequences of the acts that the theory requires: call these the theory's *direct* consequences. The second are other consequences of the procedure that the theory employs to yield its conclusions: call these its *indirect* consequences. Unless act-utilitarians manage invariably to coordinate to produce an optimum, universal satisfaction of CU has better direct consequences than universal satisfaction of AU. But satisfaction of CU may not have better indirect consequences.

AU is exclusively act-oriented. This means that it provides agents with a criterion for choosing among the acts open to them, but it does not require the use of any particular procedure in making that choice. CU is not exclusively act-oriented. It dictates a decision-making procedure; the right act is simply the one selected by that procedure. One of the indirect consequences of a theory is the expenditure of time and energy needed to apply it. No theory is costless in application, but complex and difficult decision-making procedures are more costly than simple and easy ones. The procedure required by CU is complex and difficult. AU is more flexible in allowing agents to resort to a variety of procedures and devices in response to the exigencies of actual situations. AU has direct costs since its universal satisfaction does not ensure the production of an optimum; CU has indirect costs since its procedure consumes time and energy. The direct advantages of CU must, therefore, be balanced against the indirect advantages of AU. How this balance will

turn out is an empirical question, which will depend upon the frequency with which groups of act-utilitarians achieve optimal coordination, the extent of the losses (the gap between optimum and suboptimum) on the occasions when they fail, the costs of the devices available to act-utilitarians, and the costs of the procedure required by CU. It is far from clear that the balance will, all things considered, favor CU. CU, in other words, may have no utilitarian advantage over AU.

The traditional strength of AU has been its flexibility. While it provides an objective measure for the moral value of actions, it allows agents who are attempting to satisfy it to use a wide variety of decision-making procedures that could enable them to overcome a number of common liabilities: their own fallibility, their tendency toward partiality, and their limited knowledge of the circumstances that will affect the consequences of their acts. AU provides a standard for evaluating all of these procedures: the best set of procedures, adapted to different circumstances, is the set that best satisfies AU over the long run. CU provides one such procedure, among others, for use in coordination problems. Where its direct benefits outweigh its indirect costs act-utilitarians have a reason for employing it; where its costs outweigh its benefits, they have no such reason. When both direct and indirect consequences are taken into account, CU takes its place as another strategy available to act-utilitarians who are attempting to apply their theory in an imperfect world.

* * *

A version of utilitarianism ought to be adequate by utilitarian standards. Regan has made a significant contribution to the internal evaluation of forms of utilitarianism by distinguishing clearly between their adequacy on the individual and collective levels, and by exposing the theoretical deficiency of AU on the latter level. The proper conclusions to be drawn concerning the relative merits of AU and CU are, however, more indeterminate than Regan would wish. If we confine ourselves to the direct consequences of theories, then the utilitarian costs of AU's lack of PropCOP will depend on the extent to which act-utilitarians manage in real-life situations to coordinate to produce an optimum. And when we expand attention, as utilitarians must, to all of the consequences of theories, direct and indirect, then it is far from obvious that CU is superior to AU on either the individual or collective level.

I would not wish to leave the impression that the criticisms I have made of Regan's central argument are novel. Indeed, most of them are anticipated and dealt with in the text, though I believe that they

have greater force than Regan has been prepared to concede. Be that as it may, the true measure of the value of this book will lie in the quality of the controversy that it seems likely to inspire.