



University
of Glasgow

Brown, C. (2011) Consequentialize this. *Ethics*, 121 (4). pp. 749-771.
ISSN 0014-1704

Copyright © 2011 University of Chicago Press

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

Content must not be changed in any way or reproduced in any format or medium without the formal permission of the copyright holder(s)

When referring to this work, full bibliographic details must be given

<http://eprints.gla.ac.uk/78726/>

Deposited on: 13 June 2013

Consequentialize This*

Campbell Brown

To 'consequentialize' is to take a putatively nonconsequentialist moral theory and show that it is actually just another form of consequentialism. Some have speculated that every moral theory can be consequentialized. If this were so, then consequentialism would be empty; it would have no substantive content. As I argue here, however, this is not so. Beginning with the core consequentialist commitment to 'maximizing the good', I formulate a precise definition of consequentialism and demonstrate that, given this definition, several sorts of moral theory resist consequentialization. My strategy is to decompose consequentialism into three conditions, which I call 'agent neutrality', 'no moral dilemmas', and 'dominance', and then to exhibit some moral theories which violate each of these.

I. INTRODUCTION

Arguing with a consequentialist can be frustrating. Witness a typical sort of exchange: You—a nonconsequentialist, let's assume—begin with your favorite counterexample. You describe some action, a judge's convicting an innocent man to avert a riot, say, or a doctor's murdering a healthy patient for her organs, which, so you claim, would clearly have the best consequences, yet equally clearly would be greatly immoral. So consequentialism is false, you conclude; sometimes a person ought not to do what would have best consequences.

"Not so fast," comes the consequentialist's reply. "Your story presupposes a certain account of what makes consequences better or worse, a certain 'theory of the good,' as we consequentialists like to say. Consequentialism, however, is not wedded to any such theory. We already knew that combining consequentialism with some theories of the good would have implausible results; that's what utilitarianism has

* Of the many people with whom I have had helpful discussions on the topic of this article, I remember the following: Selim Berker, Daniel Cohen, James Dreier, Andy Egan, David McCarthy, Martin Peterson, Douglas W. Portmore, Wlodek Rabinowicz, Michael Ridge, Michael Smith, and David Sobel. I thank them all, and also those I've forgotten.

taught us. In order to reconcile consequentialism with the view that this action you've described is wrong, we need only to find an appropriate theory of the good, one according to which the consequences of this action would not be best. You say you're concerned about the guy's rights? No worries; we'll just build that into your theory of the good. Then you can be a consequentialist too."

And just like that, you've been assimilated. It's a cunning ploy. Instead of showing that your nonconsequentialism is mistaken, the consequentialist shows that it's not really nonconsequentialism; instead of refuting your view, she 'consequentializes' it. If you can't beat 'em, join 'em. Better still, make 'em join you.

One may suspect, however, that this consequentializing strategy, if used in excess, will ultimately undermine consequentialism. It might succeed in immunizing consequentialism against counterexamples only at the cost of severely weakening it, perhaps to the point of utter triviality. So effortlessly is the strategy deployed that some are led to speculate that it is without theoretical limits: *every* moral view may be dressed up in consequentialist clothing.¹ Take any theory traditionally thought of as opposed to consequentialism, and it will be possible, they predict, to devise a consequentialist "counterpart theory" which mimics perfectly the verdicts of the original regarding the moral status of actions, whether they are right or wrong, permitted or forbidden.² If so, then consequentialism has something to offer everyone. No matter your ethical persuasion, there's a consequentialist theory just right for you. But then, it seems, consequentialism would be empty—trivial, vacuous, without substantive content, a mere tautology. The statement that an action is right if and only if (iff) it maximizes the good would entail nothing more substantive than the statement that an action is right iff it is right; true perhaps, but not of much use.

As I aim to show here, however, this speculation is mistaken. There are in fact limits to consequentialization. Some views of right and wrong are incompatible with consequentialism, no matter what theory of the good one adopts. These views cannot be consequentialized.

I'll proceed as follows. In Section II, I'll say more precisely what consequentialization entails. This of course will involve saying more

1. Among 'consequentializers' who make this claim, or something close to it, may be included James Dreier, "Structures of Normative Theories," *Monist* 76 (1993): 22–40; Jennie Louise, "Relativity of Value and the Consequentialist Umbrella," *Philosophical Quarterly* 54 (2004): 518–36; and Douglas W. Portmore, "Consequentializing Moral Theories," *Pacific Philosophical Quarterly* 88 (2007): 39–73.

2. This notion of a 'counterpart theory' is taken from Louise, "Relativity of Value and the Consequentialist Umbrella."

precisely what consequentialism is. My intention will not be to give the one true, uniquely correct definition of consequentialism. Indeed, I doubt there is any such thing; ‘consequentialism’ is a term of art used by philosophers to mean different things on different occasions, none of which is obviously most deserving of the name. Rather, my approach will be to adopt the “best-case scenario” for the consequentializer. I’ll try to define consequentialism in a way that makes consequentialization as easy as possible. Then, in Section III, I’ll show that, even given this favorable definition of consequentialism, the project of consequentializing all theories cannot succeed. I’ll distinguish three familiar sorts of moral theory which resist consequentialization.

II. HOW TO CONSEQUENTIALIZE

A. *Maximizing the Good*

As I say, my aim is to be as charitable to the consequentializer as I can be. My charity must, however, have limits. Were I so generous as to allow the consequentializer to mean anything at all she likes by ‘consequentialism’, then our question truly would be trivial. Thus, I impose this minimal constraint: whatever is meant by ‘consequentialism’, it must be intelligible as an elaboration of the familiar consequentialist slogan “Maximize the good.” The non-negotiable core of consequentialism, I shall assume, is the claim that an action is right, or permissible, iff it maximizes the good. This is my starting point.

But what exactly does it mean? You might say: by ‘the good’ is meant the class of all things good, the extension of the predicate ‘good’; so to maximize the good is to make as many good things as possible. But this is implausible. Suppose you can either do *x*, producing one good thing, or do *y*, producing two. Does it follow that you will maximize the good by doing *y*? Not necessarily; for it might be that the one thing is very good, yet the other two things only moderately good, so that overall the one is more good than the other two combined.

The problem with the first answer, you might say, is that it treats being good as an all-or-nothing affair, when in fact it is more a matter of degree. Not all good things are equally good; some are better than others. This suggests another answer. To maximize the good is to make things as good as possible, the best they can be. But how are we to understand the extent to which things are good when there are more things than one and they’re good to varying extents? Some procedure such as the following might seem to hold the answer. First, partition the world into a set of mutually exclusive and jointly exhaustive things (to avoid double-counting). Then assign to each thing in the partition a number representing the degree to which it is good.

Finally, add up all the numbers. This sum will represent the extent to which things are good overall.

Perhaps this procedure can be made to work. But it rests on a controversial assumption, which the consequentialist need not accept. This assumption, sometimes called ‘atomism’, is that when a whole is composed of parts, the extent of goodness of the whole must equal the sum of the extents of goodness of the parts (at least when the parts are exclusive and exhaustive). This, of course, is just what G. E. Moore warned us not to assume, and for good reason.³ Recall the earlier example: doing x will produce one good thing; doing y will produce two. But suppose now that each of the latter two things is, considered by itself, better than the former. Does it follow that you will maximize the good by doing y ? Again, not necessarily; for it might be that, although the two things are each better in isolation, when combined, they are worse.

We should resist building atomism into consequentialism. Instead, I suggest, we should understand the extent to which things are good, more holistically, as the extent to which the world as a whole is good. As David Lewis puts it, “the way things are, at its most inclusive, means the way this entire world is.”⁴ To maximize the good, then, is to make the world, the sum of all things, as good as it can be; it is to act so as to bring about the best possible world, of those which can be brought about.⁵

More precisely, I’ll assume the following general picture. When a person acts, she faces a range of alternatives, different actions she might perform. For each action, I assume, there’s a unique possible world that would obtain if the action were performed.⁶ We can call this world the *outcome* of the action. In effect, the person faces a choice between possible worlds. Each available action is the bringing about of a possible world, its outcome the possible world thereby brought about. Consequentialism then says an action is right iff its outcome is best.

B. Some Clarifications

Before moving on, three points of clarification are in order. First, consequentialism is here to be understood as a ‘criterion of right-

3. G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903), 28.

4. David Lewis, *On the Plurality of Worlds* (Malden, MA: Blackwell, 1986), 1.

5. Compare Fred Feldman, *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy* (Cambridge: Cambridge University Press, 1997), 152.

6. This assumption is sometimes called ‘counterfactual determinism’; see, e.g., Krister Bykvist, “Normative Supervenience and Consequentialism,” *Utilitas* 15 (2003): 27–49, at 30.

ness', as it's commonly put, rather than as a 'decision procedure'.⁷ That is, it need not recommend that agents aim directly at maximizing the good when deciding what to do. Plausibly, some outcomes, for example, getting to sleep, are less well achieved by aiming directly at them, and maximizing the good may be self-defeating in this way.⁸ In that case, consequentialism, as I'm thinking of it here, would recommend that agents decide what to do in some other way, for example, by adopting "rules of thumb."⁹

Second, consequentialism, as here understood, is a purely extensional thesis. It claims only that, among possible actions, certain properties, namely those of being morally right and of maximizing the good, are necessarily coinstantiated. It makes no claim of causation, determination, or explanation; it doesn't say that acts are right *because* their outcomes are best. Nor (if this amounts to a distinct claim) does it make a claim of priority; it doesn't say that the good is prior to the right.¹⁰ Sometimes it is objected that, while the extensional thesis may be all right, consequentialism puts the cart before the horse in one or other of these ways.¹¹ Since, as I've said, I want to make things as easy as possible for the consequentializer, I here understand consequentialism extensionally, so such objections may be set aside.

Finally, since the outcome of an action is here assumed to be an entire world, a person cannot in general know what the outcomes of her actions will be. Even assuming there is a fact of the matter which world would obtain should you act in a certain way, it's exceedingly unlikely that this fact would be accessible to you. Rather, in the normal case, you'll have various partial beliefs regarding which worlds might obtain conditional upon your action. In decision theory, such beliefs are represented by 'subjective probabilities' which figure in calculations of expected desirability. Thus, I'll here be considering only what is sometimes called *objective rightness*, which depends on what in fact the outcomes of actions are, and not *subjective rightness*, which depends on the variably probable outcomes of actions, as given

7. Compare Philip Pettit and Geoffrey Brennan, "Restrictive Consequentialism," *Australasian Journal of Philosophy* 64 (1986): 438–55; David O. Brink, "Utilitarian Morality and the Personal Point of View," *Journal of Philosophy* 83 (1986): 417–38.

8. Compare Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), chap. 4.

9. Compare J. J. C. Smart, "Extreme and Restricted Utilitarianism," *Philosophical Quarterly* 6 (1956): 344–54.

10. On the 'priority of the good', see John Rawls, *A Theory of Justice* (Oxford: Oxford University Press, 1971), 31.

11. See, e.g., Elizabeth Anderson, "Reasons, Attitudes, and Values: Replies to Sturgeon and Piper," *Ethics* 106 (1996): 538–54.

by the agent's subjective probabilities. This I do solely in the interest of simplicity.

C. *Theories of the Good*

As Peter Railton explains, "one has not adopted any particular morality in adopting consequentialism unless one says what the good is."¹² By itself, consequentialism is not a particular moral theory, but it becomes one, or a part of one, when combined with a theory of the good, a theory about which worlds are better or worse than which others. There are many theories of the good with which it might be combined; so there are many particular consequentialist theories. These are unified in holding that we ought to maximize the good, but divided over what the good is. Consequentialism is the collection of all these particular theories taken together, the "genus," we might say, of which they are "species."

The most famous and oldest of these species is surely utilitarianism, which results from combining consequentialism with a hedonistic theory of the good. As contemporary consequentialists are quick to point out, however, consequentialism is not limited to utilitarianism. There is, they say, a tendency to conflate the two, which must be resisted, because it makes consequentialism appear in worse shape than it really is. Once it is divorced from the crude, naive hedonistic theory of the good, many common objections evaporate. Railton, for example, gestures at a 'pluralistic' theory of the good in which "several goods are viewed as intrinsically nonmorally valuable—such as happiness, knowledge, purposeful activity, autonomy, solidarity, respect, and beauty."¹³ Such a theory, he argues, enables the consequentialist to escape the problems of 'alienation' that beset utilitarianism.

This flexibility of consequentialism opens the door to consequentialization. It raises the possibility that, by careful selection of an appropriate theory of the good, even the most adamant opponents of consequentialism may be brought under what Jenny Louise calls "the consequentialist umbrella."¹⁴ James Dreier describes the strategy: "The main strategy for 'consequentializing' any given moral theory is simple. We merely take the features of an action that the theory considers to be relevant, and build them into the consequences. For example, if a theory says that promises are not to be broken, then we

12. Peter Railton, "Alienation, Consequentialism, and Morality," *Philosophy and Public Affairs* 13 (1984): 134–71, at 148.

13. *Ibid.*, 149.

14. Louise, "Relativity of Value and the Consequentialist Umbrella"; cf. the "consequentialist vacuum cleaner" in David McNaughton and Piers Rawling, "Agent-Relativity and the Doing-Happening Distinction," *Philosophical Studies* 63 (1991): 167–85.

restate this requirement: that a promise has been broken is a bad consequence. Notice that the weighting is not yet specified. If the theory under consideration includes an absolute side constraint against promise-breaking, then we have the consequentialist version give a lexically prior negative weight to promise-breaking.”¹⁵

It will be helpful to describe the strategy more precisely. Let M be any moral theory which we aim to consequentialize. And let C be consequentialism. So a particular consequentialist theory is equivalent to some conjunction $C\&G$, where G is a theory of the good. Say that the *deontic output* of a theory P is the set of all propositions about the rightness or wrongness of actions which are entailed by P ; and say P and Q are deontically equivalent iff they have the same deontic output.¹⁶ Then to consequentialize M , we simply find some theory of the good G such that M is deontically equivalent to the particular consequentialist theory $C\&G$. We might then call $C\&G$ a *consequentialist counterpart* of M .

But why, you might ask, should we think that this process of constructing a consequentialist counterpart for M amounts to the same thing as consequentializing it? Consider an analogy. Suppose I employ you to paint my red house blue, and you then proceed to build another house which is a perfect replica of mine except it's painted blue rather than red. Then you haven't done your job. The new house may be blue, but it's not my house. You haven't done anything to my house; it's still red. Creating a blue counterpart of my house is not a way of making my house blue. Likewise, you might say, creating a consequentialist counterpart of M is not a way of making M a consequentialist theory. The counterpart theory, $C\&G$, may be consequentialist, but that's a different theory; $C\&G$ is not M . Of course, the two theories do agree on all deontic questions, about which actions are right or wrong; by hypothesis, they're deontically equivalent. But they needn't agree on questions of the good, because M need not imply G , the theory of the good which is implied by $C\&G$.

D. 'The Good' as a Theoretical Term

To solve this problem, I suggest, the consequentializer needs to adopt a certain semantic thesis, about the meaning of the term 'the good' in consequentialism. They need to understand 'the good' as a theoretical term.

By a theoretical term, I mean a term which is introduced by a theory, in this case consequentialism, as a name for something whose

15. Dreier, "Structures of Normative Theories," 23.

16. Compare the "Deontic Equivalence Thesis" in Portmore, "Consequentializing Moral Theories."

existence the theory postulates.¹⁷ As an analogy, consider so-called ‘phlogiston theory’. In the eighteenth century, chemists theorized that all combustible bodies contained a substance that was released in combustion, and they introduced ‘phlogiston’ as a name for the substance so postulated. As we all know, the theory turned out to be false. It was refuted by the discovery that in fact nothing is released by combustion; there is no such thing as phlogiston. But the theory was immune to refutation in another way: it could not have been refuted by showing that something other than phlogiston was released in combustion. Since ‘phlogiston’ is a theoretical term introduced by phlogiston theory, it names anything, if there is anything, which has all the properties attributed to phlogiston by this theory. Were there to exist anything with these properties, then that thing, by definition, would be phlogiston. (For now I ignore the possibility of multiple realizations. If it had turned out that more than one substance was given off in combustion, then phlogiston theory might also have been refuted.)

My suggestion, then, is that the consequentializer treats ‘the good’ as a theoretical term introduced by consequentialism, in the same way that ‘phlogiston’ was introduced by phlogiston theory. That is, ‘the good’ names the thing, if there is one, which has the properties attributed to the good by consequentialism. The good is thus, by definition, that which ought to be maximized.

Actually, this definition is a bit loose. Normally, to say that a person ought to do something is to say only that her doing it is necessary, though perhaps not sufficient, for her to act rightly. One ought to keep to the left side of the road when driving in Britain; but that is not to say that every motorist driving on the left—including the guy driving at twice the legal speed limit while texting on his cell phone—is acting rightly. But we wouldn’t want to define ‘the good’ as: that thing the maximization of which is necessary for acting rightly. We already know there is no such thing. Rather, there are many. For example: maximizing the number of integers. Every action maximizes this (no action results in there being fewer integers than there would have been otherwise); so, a fortiori, every right action maximizes this. And the same goes for, say, maximizing the number of prime integers; and so on. We must therefore define ‘the good’ instead as: that thing the maximization of which is both necessary and sufficient for acting rightly. That being a bit of a mouthful, however, I’ll continue using ‘that which ought to be maximized’, as a sort of shorthand.

The analogy is perhaps not perfect. In the case of phlogiston

17. Compare David Lewis, “How to Define Theoretical Terms,” *Journal of Philosophy* 67 (1970): 427–46.

theory, the term ‘phlogiston’ was not used, as we might say, ‘outside’ the theory. Proponents of the theory coined this term for their own theoretical purposes. However, the term ‘the good’ is used outside consequentialism. But this is no problem. We can think of other cases where an ordinary word is co-opted and, by stipulation, given a purely theoretical meaning. Think of ‘string’ and string theory, or, closer to home, ‘utility’ and utility theory. In treating ‘the good’ as a theoretical term, in the way I’ve suggested, we are giving it a stipulated meaning, which need not be the same as its meaning in ordinary language. Some might object that this results in an illegitimate definition of consequentialism. A proper definition, they might say, must use the term ‘the good’ as we ordinarily use it.¹⁸ If so, that would be bad news for consequentializers. But, as I’ve said, I want to make things as easy for them as I can. Hence I’ll allow them to stipulate the meaning of ‘the good’ in this way.

How does this help the consequentializer? Recall the moral theory *M* and its consequentialist counterpart *C&G*. Let ‘the good *G*’ denote whatever is the good according to *G*. As we know, an action is right according to *M* iff it is right according to *C&G*. And an action is right according to *C&G* iff it maximizes the good *G*. Therefore, an action is right according to *M* iff it maximizes the good *G*. That is, *M* implies that the good *G* (and nothing else) ought to be maximized. Given our semantic thesis, *M* thus implies that the good *G* is the good. So the consequentialization of *M* is now complete; it has been shown to be a consequentialist theory.

E. The ‘The’

You might worry that our semantic thesis has already made consequentialism a mere tautology (or at any rate, analytically equivalent to one). Consequentialism states that the good ought to be maximized, where ‘the good’ is defined as the thing which ought to be maximized. If we substitute definiens for definiendum, we thus end up with the trivial-sounding statement “The thing which ought to be maximized ought to be maximized.” However, while this may seem vacuously true, in fact it isn’t. Recall phlogiston theory, which states that phlogiston is released by combustion. Substituting definiens for definiendum as before yields: “The thing which is released by combustion is released by combustion.” This, too, may sound vacuously true. But we know it isn’t, because we know it isn’t true. And we know why: because nothing is released by combustion.

The presence of the definite article “The” is crucial here. If the

18. Compare Mark Schroeder, “Teleology, Agent-Relative Value, and ‘Good,’” *Ethics* 117 (2007): 265–95.

statement were “Anything which ought to be maximized ought to be maximized,” then it truly would be trivial; but that’s a different, non-equivalent statement. Without getting into controversial details in the philosophy of definite descriptions, we can I think safely make the following two assumptions. First, a sentence of the form “The F is G” is true iff the definite description “The F” denotes something which is G. Second, a thing is denoted by “The F” iff it, and only it, is F. It follows that a sentence of the form “The F is F” is true iff exactly one thing is F. So the sentence “The thing which ought to be maximized ought to be maximized” is true iff exactly one thing ought to be maximized.

Thus, there remains a way for a moral theory to resist consequentialization: by implying that *nothing* ought to be maximized. (And there’s another way: by implying more than one thing ought to be maximized.) As I’ll argue, some familiar moral theories do resist consequentialization in this way.

III. THE LIMITS OF CONSEQUENTIALIZATION

A. Formal Model

I begin by defining a simple formal model. The model will take entities of two sorts as primitives: possible worlds and agents. By ‘possible worlds’, I mean, in the usual way, complete ways the world might be. And by ‘agents’, I mean creatures capable of actions apt for moral evaluation. A ‘choice situation’, I’ll say, is an ordered pair $\langle i, A \rangle$, where i is an agent and A is a finite, nonempty set of worlds. (I shall briefly discuss relaxing the assumption of finiteness, made here mainly for simplicity, in Sec. III.B.3.) These worlds are the alternatives between which the agent i must choose, the outcomes of the various actions which she might perform in this situation.

I shall assume that every such ordered pair constitutes a choice situation. Admittedly, this is unrealistic. There are worlds you could never be in a position to choose, for example, worlds you don’t exist in. You might choose to cease existing, but you could not possibly choose never to have existed in the first place (by the time you make the choice, it’s already too late—you exist). There is nonetheless, I think, some sense to be made of ethical questions about such impossible choices. Would your choosing never to have existed be wrong? That question seems not entirely incomprehensible. Still, I should confess, this may be one place where I am being less charitable to the consequentializer than I could be. Suppose we were to confine attention to realistic choice situations, those where the alternatives in A are all worlds which the agent i could possibly choose. Then there may be no overlap between the alternatives in choice situations in-

volving different agents, for it could be that any world *I* could possibly choose is not one *you* could possibly choose, and vice versa. And in that case, the condition which I call below ‘Agent Neutrality’ would have no bite, for reasons that will become apparent. So I acknowledge that constraining what counts as a choice situation may be another avenue to consequentialization, though not one I will pursue here.

What I earlier called the ‘deontic output’ of a moral theory will then be represented by a ‘rightness function’ *R* which assigns to each choice situation $\langle i, A \rangle$ a subset of *A*. I shall write “ $R_i(A)$ ” to denote this subset. The worlds in $R_i(A)$ are the alternatives which are *right* (permitted) in the situation $\langle i, A \rangle$, according to the moral theory represented; the worlds in $A \setminus R_i(A)$, that is, those in *A* but not in $R_i(A)$, are the alternatives which are *wrong* (forbidden).

A theory of the good will be represented by a complete order \leq on the set of all possible worlds, or as I shall say, for short, a ‘complete world order’. Where $x \leq y$, this represents that *y* is at least as good as *x*, according to this theory of the good. By ‘order’ I mean a binary relation that satisfies two conditions: reflexivity ($x \leq x$), and transitivity (if $x \leq y$ and $y \leq z$, then $x \leq z$).¹⁹ By ‘complete order’ I mean an order which satisfies in addition a third condition: completeness (either $x \leq y$ or $y \leq x$). For now, completeness is a useful simplifying assumption. I’ll discuss relaxing it below (Secs. III.C.1 and III.D.5).

If \leq is an order on *S*, and *X* is a subset of *S*, then we may define the set of ‘greatest elements’ in *X*, relative to \leq , as the set:

$$grt_{\leq}(X) = \{x: x \in X \text{ and for all } y, \text{ if } y \in X \text{ then } y \leq x\}.$$

That is, the greatest elements are those which are at least as great as every element. (Note, I say ‘greatest elements’, not ‘maximal elements’; this distinction will be important below, in Sec. III.C.1.)

We may then say that a rightness function *R* ‘maximizes’ a world order \leq iff in every choice situation, *R* selects all and only the greatest alternatives relative to \leq ; that is, for all *i* and *A*, $R_i(A) = grt_{\leq}(A)$.

Finally, consequentialism may be defined as follows.

CONSEQUENTIALISM (C). *R* maximizes exactly one complete world order \leq .

A moral theory is consequentialistic iff it is represented by a rightness function *R* which satisfies C.²⁰

It can be shown that no rightness function maximizes more than

19. Strictly, I should say ‘preorder’.

20. This is similar to the definition of ‘teleology’ in John Broome, *Weighing Lives* (Oxford: Oxford University Press, 2004), 10–16.

one world order.²¹ So C says in effect that R maximizes at least one \leq . There are some rightness functions which do not satisfy this condition. These represent moral theories which cannot be consequentialized. Which ones are they? To give a clearer idea of this, it will be helpful to state three further conditions (for simplicity, I'll leave quantification implicit):

AGENT NEUTRALITY (AN). $R_i(A) = R_j(A)$.

NO MORAL DILEMMAS (NMD). $R_i(A) \neq \emptyset$.

DOMINANCE (D). Suppose $\{x, y\} \subseteq A \cap B$, $x \in R_i(A)$, and $y \notin R_i(A)$. Then $y \notin R_i(B)$.

C is equivalent to the conjunction of AN, NMD, and D (for proof, see the appendix, proposition 1). Thus, we may distinguish three sorts of moral theory which can't be consequentialized, corresponding to the three conditions. I'll consider each of these in turn.

B. Agent Neutrality

Utilitarianism has well-documented difficulties with people's rights. As anyone who has ever been near an ethics class will tell you, it's not hard to concoct scenarios in which achieving the greatest total happiness would require violating some poor individual's rights, for example, by denying them a fair trial, or by exploding them with dynamite to clear the opening of a cave they are obstructing. The consequentialist who rejects utilitarianism may attempt to overcome such difficulties by accounting for violations of rights in their theory of the good. They may say that worlds in which rights are violated are worse than worlds in which rights are not violated, at least other things being equal. This would be to adopt what Robert Nozick calls a "utilitarianism of rights," wherein "violations of rights (to be minimized) merely would replace the total happiness as the relevant end state in the utilitarian structure."²²

However, as Nozick points out, this approach fails to capture what many feel is a crucial aspect of the moral significance of rights, namely, that rights act as 'side constraints' against our pursuit of the good. They are such that you ought not to violate them, even in circumstances where your doing so would result in fewer violations overall.²³ Suppose that unless you kill one innocent person, thereby violating one right, someone else will kill two innocent people, thereby violating two. The view of rights as side-constraints says that you shouldn't kill the person, yet Nozick's utilitarianism of rights says you

21. Suppose $x \leq_1 y$ and $\sim x \leq_2 y$. Then R maximizes \leq_1 only if $y \in R_i(\{x, y\})$, and R maximizes \leq_2 only if $y \notin R_i(\{x, y\})$. So R cannot maximize both \leq_1 and \leq_2 .

22. Robert Nozick, *Anarchy, State, and Utopia* (Oxford: Blackwell, 1974), 28.

23. *Ibid.*, 29.

should, since this would minimize the total number of rights violations overall.

This side-constraint rights theory is an example of what are sometimes called agent-relative, or agent-centered, theories.²⁴ As usually defined, a theory is agent-relative iff it gives different aims to different agents; otherwise it's agent-neutral. Thus, for example, the side-constraint rights theory gives me the aim of seeing to it that *I* don't violate any rights, but gives you the aim of seeing to it that *you* don't violate rights. For another example, some moral theories say that agents have a special obligation to promote the well-being of their own family and friends, over and above the obligation they have to promote the well-being of people generally.²⁵ In contrast, utilitarianism is agent-neutral, because it gives everyone the same aim: maximizing total happiness.

Unsurprisingly, if R represents an agent-relative theory, then it violates the condition AN. Suppose R represents the side-constraint rights theory. And let the worlds w_m and w_y be such that in w_m you violate some right and I violate none, but in w_y I violate some right and you violate none. (Surely there are such worlds.) Since the theory says I ought to choose w_m over w_y , and you ought to choose w_y over w_m , we have $R_{me}(\{w_m, w_y\}) = \{w_m\}$ and $R_{you}(\{w_m, w_y\}) = \{w_y\}$. And this plainly violates AN, because by hypothesis $w_m \neq w_y$. So this theory resists consequentialization, as do all agent-relative theories.

1. *Agent-relative value.*—This result might be thought to reveal a certain limitation of our formal model. A popular strategy for consequentializing agent-relative theories involves devising agent-relative theories of the good, theories according to which what's good relative to one person need not be good relative to another. But the model doesn't allow the good to be agent-relative in this way.

It might seem that such a strategy is incompatible with the idea of maximizing *the* good. On this approach, one might say, there is not a single thing—the good—which every agent ought to maximize; rather there are many things—the good relative to me, the good relative to you, and so on—and each agent ought to maximize a different one. Perhaps, however, this can be thought of in another way. What differs from person to person, one might say, is not what the person ought to maximize, but rather how the person maximizes it; we all ought to maximize the same thing—the good relative to the agent—but we each do so in our own way. Here's an analogy.²⁶ Suppose you and I are competing in a race. Then, in one manner of

24. See, e.g., Parfit, *Reasons and Persons*, 27.

25. See, e.g., Michael Smith, "Immodest Consequentialism and Character," *Utilitas* 13 (2001): 173–94.

26. As I recall, this example was first suggested to me by Michael Smith.

speaking, you and I have the same goal: winning the race. But what would make it the case that I achieve this goal, that is, my crossing the finish line before you, is not what would make it the case that you achieve it, that is, your crossing before me. There is a single thing, winning the race, which we both aim to achieve, but we each achieve it (if we do) in our own way: you by your winning, and me by my winning. Similarly, one might say, in the case of the rights theory, there is a single thing, respecting rights, which we both ought to maximize, but we each maximize it (if we do) in our own way: you by your respecting rights, and me by my respecting rights.

How might this sort of view be represented formally? The obvious approach, I think, is to borrow a familiar formal device for representing what is sometimes called ‘indexical content’, namely, the device of centered worlds.²⁷ Suppose you and I each utter the sentence “I have a beard.” Some hold the view that our utterances express a common content. But that content cannot be a proposition, at least not in the sense of a set of possible worlds, because the worlds in which you have a beard do not perfectly overlap those in which I do. Rather, on this approach, the content is a set of centered worlds: the set of every world-individual pair $\langle w, i \rangle$ such that i has a beard in w .

Thus, we might represent an agent-relative theory of good, not by an order of worlds (as above), but rather by an order of centered worlds. For example, a theory that values respecting rights (and nothing else) would be represented by an order \leq such that $\langle x, i \rangle \leq \langle y, j \rangle$ iff the number of rights violated by i in x is no less than the number violated by j in y . Then the greatest elements in a set of worlds A would be defined, now doubly indexed to an order \leq and an agent i , as follows:

$$grt_{\leq, i}(X) = \{x: x \in X \text{ and for all } y, \text{ if } y \in X \text{ then } \langle y, i \rangle \leq \langle x, i \rangle\}.$$

A rightness function R would be defined as maximizing a world order \leq iff, for all i and A , $R_i(A) = grt_{\leq, i}(A)$. Finally, consequentialism would say, as before, that R maximizes exactly one \leq .

This would get us part of what we wanted: some rightness functions which violate AN nonetheless maximize centered world orders. The problem, however, is that if R maximizes one, then it maximizes many; it cannot maximize exactly one. Thus, our reformulated consequentialism, far from being trivially true, is trivially false, self-contradictory. (For proof, see the appendix, proposition 2.) Roughly, the reason is this. If the centered worlds $\langle x, i \rangle$ and $\langle y, j \rangle$ are centered on different people (i.e., if $i \neq j$), then they make no difference to

27. See, e.g., David Lewis, “Attitudes De Dicto and De Se,” *Philosophical Review* 88 (1979): 513–43.

whether R maximizes \leq . That is, if \leq' is the same as \leq except for the relative positions of these centered worlds, then R maximizes \leq iff it maximizes \leq' . So R cannot maximize only one order.

This way of trying to accommodate agent-relative theories within consequentialism therefore fails. To be clear, I do not claim that agent-relative views are incoherent. I claim only that they are incompatible with the view that there is a single thing, the good, that we all ought to maximize. Such views are better thought of as the saying that there are many things that ought to be maximized, a different one for each of us.

C. No Moral Dilemmas

By a *moral dilemma*, I mean a situation in which a person cannot avoid acting wrongly. In our formal model, a choice situation $\langle i, A \rangle$ constitutes such a dilemma, according to a rightness function R , iff $R_i(A)$ is the empty set. Moral dilemmas have been widely discussed by philosophers, some of whom believe they are possible. Moreover, the possibility of moral dilemmas seems to be entailed by familiar moral theories, which might be called *absolutist theories*, since they incorporate absolute prohibitions. Consider, for example, a theory which holds that violations of rights are absolutely morally forbidden; it is always wrong in any possible situation to violate a right. Suppose, further, that the catalog of rights endorsed by this theory is such that sometimes a person cannot help but violate at least one right. Then this theory cannot be represented by a rightness function which satisfies NMD, and so it cannot be consequentialized.

1. *Going incomplete*.—Again, this might be thought to show a limitation of our formal model. In our model, the good must be represented by a *complete* order. In effect, this rules out *incommensurability*. That is, it implies that for any two worlds, either one is better than the other, or they are equally good. Suppose, then, we were to relax this requirement and allow the good to be represented by an incomplete, or partial, order. If \leq is incomplete, then there is at least one set of worlds A with *no* greatest elements relative to \leq ; that is, $grt_{\leq}(A) = \emptyset$. Thus, a rightness function which violates NMD, though it cannot maximize any *complete* order, might nonetheless maximize an incomplete one (and, as before, it cannot maximize more than one). Dropping the assumption of completeness would therefore allow moral dilemmas within consequentialism.

However, if we allow the good to be incomplete in this way, then we should, I think, revise our definition of maximization. We should define this, not in terms of greatest elements, as above, but rather in terms of maximal elements. The maximal elements in a set X , relative to an order \leq , may be defined as follows:

$$\max_{\leq}(X) = \{x: x \in X \text{ and for all } y, \\ \text{if } y \in X, \text{ and } x \leq y \text{ then } y \leq x\}.$$

(See Sec. III.A for the definition of greatest elements.) Thus, an element is maximal iff it is at least as great as every element that is at least as great as it. If \leq is complete, then $\max_{\leq}(X) = \text{grt}_{\leq}(X)$, and so it didn't matter before whether maximization was defined in terms of greatest or maximal elements. But if \leq is incomplete, then $\text{grt}_{\leq}(X)$ may be a proper subset of $\max_{\leq}(X)$. Every greatest element must be maximal, but the converse need not hold. So long as X is finite, $\max_{\leq}(X)$ is nonempty. Defining maximization in terms of maximal elements, as I think we should, would therefore again make consequentialism incompatible with moral dilemmas, regardless of whether the good is complete.²⁸

Here is my argument for defining maximization in terms of maximal elements. Suppose that, having promised to show you "one of the most impressive sights in Great Britain," I take you to see Salisbury Cathedral.²⁹ Were you later to learn of some more impressive sight in Britain, you would then be entitled to complain that I had not fulfilled my promise. (There are of course contexts in which one can say of a sight which is only, say, the fifth most impressive in Britain that it is "one of the most impressive." But that would be speaking a little loosely, like saying, e.g., that someone is "at the front of the queue" when in fact there are several people in front of her. Assume then that when I gave my promise, I made it clear that I was speaking quite strictly; e.g., I said that the sight I would show you was "one of the *very most* impressive," that it was "unsurpassed" or "second to none" in its impressiveness.) Suppose what you learn, however, is only that, though no sight in Britain is more impressive than Salisbury Cathedral, there is one sight, Stonehenge, such that Salisbury Cathedral is not at least as impressive as it. Then surely you have no such complaint. You might ask: "If Salisbury Cathedral is not at least as impressive as Stonehenge, then why did you not take me to see Stonehenge instead?" But then I could easily answer this challenge by saying: "Because Stonehenge is not more impressive than Salisbury Cathedral."

This suggests the following general principle: if an F is not one of the most G F s, then some F is more G than it. (For example, substituting "puppy" and "cute" for " F " and " G ," respectively, yields this

28. The definition would be this: R maximizes \leq iff $R_i(A) = \max_{\leq}(A)$ for all i and A .

29. This example was inspired by John Broome, "Is Incommensurability Vagueness?" in *Ethics out of Economics* (Cambridge: Cambridge University Press, 1999), 123–44.

instance of the principle: if a puppy is not one of the most cute puppies, then some puppy is more cute than it.) And this implies that in any choice situation, any alternative that is not worse than any other alternative is one of the best alternatives. But an alternative that is not less great than any other, relative to an incomplete order, might nonetheless fail to be a greatest alternative relative to that order (though it must be maximal; recall we're assuming A must be finite). Thus, if maximization were defined in terms of greatest elements, consequentialism would allow that an alternative might be wrong even though it is one of the best. And that, it seems to me, is something that consequentialism cannot allow. Consequentialism must say at least this much: if an alternative is one of the best, then it is permissible.

Interestingly, however, if we allow incompleteness and define maximization in terms of maximal elements, then, though consequentialism would still imply NMD, it would fail to imply another of the three conditions, namely, D. I shall return to this in Section III.D.5.

2. *Infinitely many alternatives.*—Another way to accommodate moral dilemmas within consequentialism would be to drop the assumption that all choice situations involve only finitely many alternatives. If A is infinite, then even $\max_{\leq}(A)$ may be empty. However, moral theories which imply dilemmas in finite cases would still elude consequentialization.

D. Dominance

Condition D may be the least intuitive of the three. It requires the following. Suppose that in a given choice situation, two worlds x and y are among the alternatives. And suppose in this situation, x is right and y wrong. Then x dominates y in the following sense: y cannot be right in any situation where x is an alternative; the presence of x is always sufficient to make y wrong.

I'll discuss three sorts of D-violating moral theory.

1. *Satisficing.*—One place to look for such theories is with those that relax the maximizing element in consequentialism. To some, it has seemed that accepting nothing less than the best, as consequentialists do, is too demanding. They believe it's sometimes okay to do something which is worse than the best, provided it's not too much worse. Theories which are less demanding in this way are sometimes called *satisficing theories*, the idea being that 'satisficing' is a more tolerant version of maximizing.³⁰

It will be useful to discuss a particular (made up) satisficing the-

30. See, e.g., Michael Slote, *Common Sense Morality and Consequentialism* (London: Routledge, 1985).

ory. Consider then ‘satisficing utilitarianism’, defined as follows: an alternative x is right iff the total well-being in x is at least as great as 0.8 times the total well-being in every alternative. Suppose there are two choice situations as shown in the following table.

Situation 1		Situation 2	
Worlds	Well-being	Worlds	Well-being
w_1	100	w_1	100
w_2	80	w_2	80
w_3	120	w_4	60

In each situation, there are three alternatives, two of which— w_1 and w_2 —are available in both. The difference is that w_3 , which is available in the first situation, is replaced by w_4 in the second. The numbers in the “Well-being” columns represent the total well-being in each of the worlds. Suppose R represents satisficing utilitarianism. Then we have:

$$R_i(\{w_1, w_2, w_3\}) = \{w_1, w_3\}$$

$$R_i(\{w_1, w_2, w_4\}) = \{w_1, w_2\}.$$

But then R cannot satisfy D. If it did, then since w_2 is wrong in the first situation, where w_1 is present and not wrong, w_2 would be dominated by w_1 , and so it couldn’t be right in the second situation, where w_1 is also present. Thus, satisficing utilitarianism cannot be consequentialized.

According to satisficing theories, then, there is no such thing as the good. This might seem a puzzling conclusion. Think of the way we would normally describe the difference between satisficing utilitarianism and regular utilitarianism. These two theories, we would say, share a common theory of the good: they both say that the goodness of a world is given by the total well-being it contains. Their difference is over how we should “respond” to the good: roughly, whereas regular utilitarianism says we must maximize the good, satisficing utilitarianism says it’s okay if we merely satisfice it. However, our conclusion entails quite a different picture. The two theories disagree about the good: whereas regular utilitarianism says there is such a thing as the good, satisficing utilitarianism says there is no such thing.

But the puzzle is easily resolved if we bear in mind the definition of ‘the good’ we’ve been assuming. On this definition, ‘the good’ names the thing, if there is one, which ought to be maximized. Obviously, when proponents of satisficing theories say something like “It is sometimes permissible not to maximize the good,” they cannot

mean ‘the good’ in *this* sense; that would be incoherent. They must mean it in some other sense.

2. *Self and other*.—A second theory which violates D is so-called ‘self-other utilitarianism’, as proposed by Ted Sider.³¹ The rationale for this view is to allow people to sacrifice their own well-being in certain situations where this wouldn’t be allowed by ordinary utilitarianism. It may be defined as follows: an alternative x is right iff either (1) the total well-being of all people is at least as great in x as in every alternative, or (2) the total well-being of all people excluding the agent, i , is at least as great in x as in every alternative.

To see how this violates D, consider three worlds that contain only two people, i and j , as given in the following table:

Worlds	Well-being	
	i	j
w_1	1	1
w_2	0	1
w_3	1	0

Suppose first that i must choose between all three worlds. Then self-other utilitarianism says i may choose either w_1 , which would maximize the total well-being of everyone, or w_2 , which would maximize the total well-being of everyone excluding i , but not w_3 , which would maximize neither. So, in the sense introduced above, w_2 dominates w_3 ; w_2 is right in a situation where w_3 is wrong. But now suppose i has only w_2 and w_3 to choose from. Now self-other utilitarianism says both alternatives are right, thereby violating dominance.³²

3. *Pareto*.—My final example of a D-violating theory has not, to my knowledge, been explicitly proposed by anyone, but it seems in the neighborhood of views one sometimes finds among economists.

Suppose that, though you are sympathetic to the general utilitarian approach, you ultimately reject utilitarianism on the specific grounds that interpersonal comparisons of well-being are meaningless. Nonetheless, you think utilitarianism gets things right in those exceptional cases where interpersonal comparisons are not required, namely, those cases where one alternative is ‘Pareto-superior’ to the rest (i.e., the alternative would make at least one person better off, and no person worse off, than any other alternative). Then you might be tempted by the following view, which I’ll call ‘Paretian utilitarianism’: (1) an

31. Theodore Sider, “Asymmetry and Self-Sacrifice,” *Philosophical Studies* 70 (1993): 117–32.

32. Douglas W. Portmore has advanced a variant of this sort of view which he calls “Schefflerian utilitarianism”; see Portmore, “Consequentializing Moral Theories,” 57. This view, as I understand it, similarly violates D.

alternative is wrong if it is Pareto-inferior to some other alternative; (2) otherwise it is right. (Note: this differs from ordinary utilitarianism only in the second clause. Utilitarianism allows that an alternative may be wrong even when not Pareto-inferior.)

Now consider these alternatives:

Worlds	Well-being	
	<i>i</i>	<i>j</i>
w_1	2	0
w_2	0	1
w_3	1	0

Suppose first that all three worlds are alternatives. Then Paretian utilitarianism rules out w_3 , because it is Pareto-inferior to w_1 , but it permits both w_1 and w_2 . (Perhaps you are tempted to say that w_2 is also clearly worse than w_1 , because it's better to give one person a benefit of two than to give another person a benefit of one. But that argument rests on an interpersonal comparison of well-being, which we're assuming is impossible.) So w_2 should dominate w_3 , according to D. But then, when the choice is just between w_2 and w_3 , Paretian utilitarianism permits both, thereby violating D.

4. *Relative value again.*—One might try to accommodate such theories within consequentialism by introducing another sort of relative goodness, though this time relativized to sets of alternatives, rather than to agents. For example, in the case of satisficing utilitarianism discussed above, one might say that whether w_1 is better than w_2 depends on what the other alternatives are: in situation 1, where the only other alternative is w_3 , w_1 is better than w_2 ; but in situation 2, where the only other alternative is w_4 , it isn't. Dominance depends on the assumption that the value of an alternative remains constant between different choice situations. Thus, were we to reject that assumption, we would have to reject D too.

However, since this proposal involves relative goodness, it is vulnerable to the same objection that I made above against agent-relative goodness (Sec. III.B.1). In brief: the most natural way to represent it formally would be to represent the good by an order of centered worlds, though this time centered on sets of alternatives, rather than agents; but that would make consequentialism trivially false, for the reasons I gave above.

5. *Incompleteness again.*—I argued above (Sec. III.C.1) that if we drop the assumption of completeness, we should define maximization in terms of maximal elements. And I pointed out that, if we did this, then consequentialism would no longer imply D. However, this would

not help with consequentializing the D-violating theories discussed above.

True, each of these theories does maximize (in this sense) at least one incomplete world order. For example, satisficing utilitarianism maximizes the order \leq defined as follows: for all worlds x and y , $x \leq y$ iff either (a) $x = y$, or (b) the total well-being in x is less than 0.8 times the total well-being in y . However, they also all maximize more than one incomplete order, given the very innocuous assumption that worlds may differ without differing with respect to people's well-being. For example, suppose that $w_1 \neq w_2$ and that total well-being is the same in w_1 and w_2 . Then define another order \leq' that is identical to the one above with the exception that $w_1 \leq' w_2$ and $w_2 \leq' w_1$. This order is also maximized by satisficing utilitarianism. According to \leq' , w_1 and w_2 are equally good, whereas according to \leq , they are incommensurable. Thus, the problem, we might say, is that satisficing utilitarianism gives us no way to distinguish equality and incommensurability, in some cases.

Indeed, if consequentialism is defined in this way—that is, allowing incompleteness, with maximization defined in terms of maximal elements—then it implies that there is no equality except in the trivial sense that every world is equal in value with itself. (For proof, see the appendix, proposition 3.) This suggests to me that this is not an attractive way to define consequentialism.

IV. CONCLUDING REMARKS

So the project of consequentializing every ethical theory fails. But what does this show? It might seem unclear whether this failure is good or bad news for consequentialism. In my view, it's good.

I'm most confident of this in the case of moral dilemmas. It seems to me a substantive issue whether moral dilemmas are possible. That consequentialism takes a side on this issue, denying their possibility, shows therefore that it has substantive content; it's not empty or vacuous. Moreover, it seems to me, consequentialism takes the *right* side on this issue. I find moral dilemmas implausible. So this is good news for consequentialism. It would be worse for consequentialism if it either took the wrong side, or took no side at all. I'm less confident that consequentialism is on the right side with respect to other issues discussed above, for example, agent-relativity. Still, it's a good feature of consequentialism, at least *pro tanto*, that it has substantive implications in these areas.

I'll close by drawing out another moral of my conclusion, related to something Dreier says. Dreier's motivation for consequentializing is that he wants to overcome a certain "stigma" which he says afflicts de-

fenders of “common sense morality” when they try to deny consequentialism.³³ To deny consequentialism, he says, they must claim that we are sometimes required to do less good than we might, but that claim has a “paradoxical air.”³⁴ So defenders of commonsense morality, who deny consequentialism, are stigmatized as having a seemingly paradoxical position. (Note the semantic thesis adopted above provides a nice explanation of the air of paradox. If the good is, by definition, the thing which ought to be maximized, then to say that a person ought not to maximize the good is to say that they ought not to maximize the thing they ought to maximize, and that’s obviously paradoxical.)

Dreier thinks the way to avoid the stigma is to avoid denying consequentialism. If we consequentialize commonsense morality, then defenders of commonsense morality need not deny consequentialism. If I’m right, however, this way of avoiding the stigma doesn’t work; some elements of commonsense morality—that is, agent-centeredness, moral dilemmas, satisficing—prevent its consequentialization. But there’s a better way to avoid the stigma. Nonconsequentialists need not accept the paradoxical claim that sometimes we ought not to maximize the good. Instead, they can claim that there is no such thing as the good.

Appendix

PROPOSITION 1: R satisfies C iff R satisfies AN, NMD, and D.

Proof. The left-to-right implication is obvious. We prove only the converse. Suppose R satisfies AN, NMD, and D. We then need only to show that R maximizes at least one complete world order, because, as shown in footnote 21, it cannot maximize more than one.

Let i be any agent. And define \leq by: $x \leq y \Leftrightarrow y \in R_i(\{x, y\})$. We need to show that \leq is a complete order (i.e., transitive and complete), and that $R_j(A) = grt_{\leq}(A)$ for all j and A .

Completeness follows immediately from NMD. To prove transitivity, assume for reductio: $x \leq y$, $y \leq z$, and $\sim x \leq z$. Since $\sim x \leq z$, the definition of \leq implies $z \notin R_i(\{x, z\})$, and then NMD implies $R_i(\{x, z\}) = \{x\}$. It follows by D that $z \notin R_i(\{x, y, z\})$. But since $y \leq z$, $z \in R_i(\{y, z\})$, and so by D, $y \notin R_i(\{x, y, z\})$. And since $x \leq y$, it follows by parallel reasoning that $x \notin R_i(\{x, y, z\})$. But by definition, $R_i(\{x, y, z\}) \subseteq \{x, y, z\}$. So $R_i(\{x, y, z\}) = \emptyset$, and therefore R doesn’t satisfy NMD, violating our assumption.

Now we prove $R_j(A) = grt_{\leq}(A)$. Since R satisfies AN, it will be sufficient to prove $R_i(A) = grt_{\leq}(A)$. Assume first $x \in R_i(A)$, it follows

33. Dreier, “Structures of Normative Theories,” 24.

34. *Ibid.*

from D and NMD that, for all $y \in A$, $x \in R_i(\{x, y\})$. So, by the definition of \leq , $x \in \text{grt}_{\leq}(A)$. Next assume $x \notin R_i(A)$. If $x \notin A$, then $x \notin \text{grt}_{\leq}(A)$. If $x \in A$, then, by NMD and D, there exists some $y \in A$ such that $x \neq y$ and $x \notin R_i(\{x, y\})$, in which case $x \notin \text{grt}_{\leq}(A)$. Therefore, $R_i(A) = \text{grt}_{\leq}(A)$.

PROPOSITION 2: Suppose \leq is an order of centered worlds and that R maximizes \leq ; that is, $R_i(A) = \text{grt}_{\leq, i}(A)$ for all i and A . Then there exists some centered world order \leq' such that $\leq \neq \leq'$ and R maximizes \leq' .

Without loss of generality, assume there are only two agents, a and b . Then define centered world orders \leq^1 and \leq^2 as follows:

- if $i = j$, $\langle x, i \rangle \leq \langle y, j \rangle$ iff $\langle x, i \rangle \leq \langle y, j \rangle$ iff $\langle x, i \rangle \leq_2 \langle y, j \rangle$
- if $i = a$ and $j = b$, $\langle x, i \rangle \leq_1 \langle y, j \rangle$ and $\sim \langle x, i \rangle \leq_2 \langle y, j \rangle$
- if $i = b$ and $j = a$, $\sim \langle x, i \rangle \leq_1 \langle y, j \rangle$ and $\langle x, i \rangle \leq_2 \langle y, j \rangle$

It is easily seen that if R maximizes \leq , then R maximizes both \leq_1 and \leq_2 . But by construction $\leq_1 \neq \leq_2$, so either $\leq \neq \leq_1$ or $\leq \neq \leq_2$.

PROPOSITION 3: Suppose \leq is the only world order (complete or incomplete) that is maximized by R , in the sense that $R_i(A) = \max_{\leq}(A)$ for all i and A . Then, for all worlds x and y , if $x \neq y$, either $\sim x \leq y$ or $\sim y \leq x$.

Define \leq' as follows: $x \leq' y$ iff either (a) $x = y$ or (b) $x \leq y$ and $\sim y \leq x$. It follows that $\max_{\leq}(A) = \max_{\leq'}(A)$ for all A , and so if R maximizes \leq , it maximizes \leq' . But, by hypothesis, \leq is the only order maximized by R , and so $\leq = \leq'$. And, by definition, for all x and y , if $x \neq y$, either $\sim x \leq' y$ or $\sim y \leq' x$.