

THE TELEOLOGICAL/DEONTOLOGICAL DISTINCTION

PETER VALLENTYNE

Department of Philosophy, University of Western Ontario, London, Ont., Canada N6A 3K7

1. Introduction

The teleological/deontological distinction was introduced in 1930 by C.D. Broad,¹ and since then it has come to be accepted as *the* fundamental classificatory distinction for moral philosophy. I shall argue that the presupposition that there is a *single* fundamental classificatory distinction is false. There are too many features of moral theories that matter for that to be so. I shall argue furthermore that as it is usually drawn the teleological/deontological distinction is not even *a* fundamental distinction. Another distinction, that between theories that make the right depend solely on considerations of goodness (axiological theories) and those that do not, is significantly more important.

2. Deontological theories

Act utilitarianism is a paradigm teleological theory, and The Divine Command Theory and Kant's moral theory are paradigm deontological theories. It is unclear, however, which of the many features of these theories are the defining characteristics of teleological and deontological theories respectively. In order to assess the importance of the teleological/deontological distinction, we need first to clarify the nature of the distinction. Let us start therefore by examining some of the different characterizations that have been given.

Often authors (e.g., Rawls²) characterize deontological theories simply as theories that are not teleological. So characterized, the nature of deontological theories depends on how teleological theories are characterized. In a later section we shall examine some of the characterizations of teleological theories. In this section we shall examine some of the characterizations of deontological theories that are independent of the characterization of teleological theories.

It is often claimed that deontological theories – but not teleological theories – are *rule-based*, i.e., assess the permissibility of actions in terms of whether they conform to some specified set of rules. The problem with this characterization of deontological theories is that *all* theories other than those that merely specify *prima facie* considerations are rule-based. Act utilitarianism, for example, is rule-based, since it assess the permissibility of actions in terms of whether they conform to the rule “Maximize the goodness of consequences!”. Both deontological

and teleological theories may be (and generally are) rule-based. The difference between the two lies in the *kinds* of rules that they invoke.

In a similar vein it is sometimes claimed that deontological theories, or at least one type thereof, are *absolutist* in that they claim that there are certain *kinds* of actions that are *absolutely* obligatory or forbidden (i.e., such that all actions of that type are obligatory or forbidden respectively.)³

There are two problems with this claim. First of all, without some restriction on the admissible action-types *all* theories are absolutist in this sense, since all theories hold that all actions of the type “is permissible” are permissible. Secondly, even if an appropriate restriction is placed on the admissible action-types to avoid this trivialization (perhaps allowing only non-normative action types), the claim is still false. Act utilitarianism, a paradigm non-deontological theory, is absolutist in this sense. This is because any action that falls under the description “does not have consequences that contain as much happiness as is feasible” is judged forbidden – no matter what its other characteristics are. Furthermore, Ross’s intuitionist theory, a paradigm deontological theory, is non-absolutist, because it provides neither necessary nor sufficient conditions for the permissibility of action, but only *prima facie* considerations that must somehow be weighed together. The characterization of deontological theories as absolutist is thus inadequate.

A closely related characterization of deontological theories is as theories for which there are certain kinds of actions that are always (or never) forbidden, where whether or not a given action is of the specified kinds *does not depend on what its outcome is*.⁴

The distinction between theories that base the permissibility of actions on their outcomes and those that do not is an important distinction. The outcome of an action is a state of affairs that would be realized if the action were performed. It is objectively determined in that its determination is independent of what the agent, or anyone else, believes (e.g., about what would happen if the action were performed). Theories that do not base the permissibility of actions on their outcomes generally base it on their intended, anticipated, or reasonably anticipatable outcomes. Unlike (objective) outcomes, neither intended nor anticipated outcomes are objectively determined, since what they are depends only on the agent’s mental state (what he/she intends or believes), and not on what would happen if the action were performed. And the reasonably anticipatable outcome is neither objectively determined, nor determined on the basis of the agent’s mental state, but rather on the basis what it would be (intersubjectively) reasonable to anticipate happening.

Characterizing deontological theories as theories for which the permissibility of actions does not depend on their outcomes, however, is intuitively inadequate. The Ten Commandment Theory – a paradigm deontological theory, according to which an action is permissible just in case it conforms to the Ten Commandments – contains the injunction “Do not kill!”, and therefore grounds the permissibility of actions in their outcomes. Whether or not an action is a killing depends on what its outcome is, and not, for example, on what its anticipated or intended outcome is. The above characterization mistakenly classifies The Ten Commandment Theory

as non-deontological. Furthermore, a form of utilitarianism that grounds the permissibility of actions in the goodness of their anticipated outcomes, a paradigm non-deontological theory, does not ground the permissibility of actions in features of their outcomes, but rather in features of the agent's mental state. The above characterization of deontological theories mistakenly classifies it as deontological.

Another closely related characterization of deontological theories is as theories that ground the permissibility of actions in and only in their *intrinsic nature*.⁵ The problem here is that it is not clear what it might mean to say that a particular feature is part of the "intrinsic nature" of an action.

One possibility is that the intrinsic nature of an action consists of its "relatively immediate" features.⁶ There are various sorts of reasonably intuitive ways of interpreting 'immediate'. One notion of immediacy that has received a fair amount of attention recently is that implicit in the distinction between what the agent "brings about" and what he/she merely "allows to happen".⁷ Roughly speaking an agent brings about a state of affairs just in case it would not have been realized, if the agent had been passive, i.e., if he/she had not intervened into the "normal course of events". For example, under normal circumstances, if an agent pushes someone off a high cliff and the person dies, then his/her death is something that the agent brings about. On the other hand, if the agent stands by and "does nothing" as the person falls off the edge of the cliff, his/her death is something that the agent merely allows to happen.

Although this is but one notion of immediacy, it seems clear that there is something to the distinction between theories that ground the permissibility of actions solely in their relatively immediate features and those that do not. This distinction does not, however, capture the intuitive deontological/non-deontological distinction. A theory that judges an action permissible just in case its immediate outcome (however construed) is maximally good grounds the permissibility of actions solely in their immediate features, but intuitively is not a deontological theory. Furthermore, a theory that consists of the injunction "Don't bring about or allow the extinction of any species!" does not ground the permissibility of actions solely in their immediate features (since what one allows to happen in the distant future is presumably not an immediate feature), but intuitively is a deontological theory. Similar counterexamples could, I suggest, be found for other interpretations of immediacy.

Another possibility is that the intrinsic nature of an action consists of its *non-comparative* features, i.e., those features that it has independently of the features that its alternatives have. On this interpretation the feature of being a killing is the type of feature that can be part of the intrinsic nature of an action, but the feature of having maximally good consequences is not. So interpreted, however, the characterization of deontological theories (or even one type thereof) as theories that ground the permissibility of actions in and only in their intrinsic nature is intuitively inadequate. Minimal act utilitarianism – according to which an action is permissible just in case its consequences are not bad (e.g., the net level of happiness is not negative) – grounds the permissibility of an action in and only in its non-comparative features, but intuitively it is not deontological. And a theory whose sole injunction is "Don't kill unless it is necessary to save the life of a close friend

or family member!”, does ground the permissibility of an action in features of its alternatives (on whether there is an alternative that saves the life of a close friend or family member, without killing anyone else), but intuitively is deontological.

So far, then, we have judged the following characterizations of deontological theories to be inadequate: (1) as rule-based, (2) as absolutist, (3) as not grounding the permissibility of actions in their outcomes, (4) as grounding the permissibility of actions in and only in their intrinsic nature, where the intrinsic nature of action is understood as (a) its “relatively immediate” features, or (b) its non-comparative features.

A feature common to all of these characterizations of deontological theories is that they do not impose any requirement concerning how the good and the right are related. Let us now examine some characterizations that do impose some such requirement.

Sometimes, deontological theories are characterized as theories for which the right is prior to the good, i.e., for which the good depends on the right.⁸ This characterization, however, is too narrow to be adequate. Theories for which the right is prior to the good (such as Kant’s) are indeed intuitively deontological, but they are not the only theories that are deontological. There are at least two kinds of theories for which the right is not prior to the good which are intuitively deontological. (1) Theories that make the right and the good *independent* of each other are generally considered to be deontological. For example, a theory that assesses the goodness of states of affairs in terms of the total amount of happiness they contain, and that judges an action permissible just in case it conforms to The Ten Commandments is intuitively deontological. It does not, however, make the right prior to the good. (2) Theories for which the right is not prior to the good, but rather depends on considerations of goodness, but does not depend *solely* on such considerations are generally considered to be deontological. For example, a theory that assesses the goodness of states of affairs in terms of the total amount of happiness they contain, and which directs the agent to produce the best consequences subject to a constraint against killing is intuitively deontological. It does not, however, make the right prior to the good.

A more adequate characterization of deontological theories is as theories that do not make the right depend *solely* on considerations of goodness.⁹ None of the intuitive paradigms of deontological theories (e.g., Kant’s theory, The Divine Command Theory, The Ten Commandment Theory, and Ross’s non-absolutistic theory) ground the permissibility of actions solely in considerations of goodness. And all of the intuitive paradigms of non-deontological theories (e.g., forms of utilitarianism based on the goodness of outcomes, anticipated outcomes, or reasonably anticipatable outcomes, and minimal act utilitarianism) do ground the permissibility of actions solely in considerations of goodness.

The distinction between theories that ground the permissibility of actions solely in considerations of goodness and those that do not is an important distinction. Indeed, I shall argue that it is more important than any of the traditional teleological/non-teleological distinctions. First, however, let us take a closer look at the feature of making the right depend solely on consideration of goodness. Following Michael Stocker¹⁰ let us call such theories *axiological*.

3. Axiological theories

A theory is *axiological* roughly just in case it makes the permissibility (and obligatoriness) of some kind of entity (e.g., actions, or social institutions) depend solely on considerations of goodness. It is of course possible for a theory to assess the permissibility of one type of entity (e.g., social institutions) solely in terms of goodness, but to assess that of another kind of entity (e.g., actions) on some other basis. In what follows we shall be concerned with the permissibility of actions only, and not that of other kinds of objects. Thus, ‘the right’ is to be understood as ‘the rightness of actions’, and ‘axiological’ as ‘act-axiological’.

Surprisingly little attention has been paid to the feature of making the right depend solely on considerations of goodness. Authors often discuss the feature of making the good prior to the right, but this is not exactly the same feature. To say that the good is prior to the right is to say that the right depends *at least in part* on considerations of goodness. It is not, however, to say that the right depends *solely* on such considerations. Consider, for example, a theory that judges an action permissible just in case *of those actions that are not killings* it has the best consequences. This theory holds that the good is prior to the right, but it does not hold that the right depends only on the good, since it makes the right also depend on which of the feasible actions are killings.

There are, then, two ways for a theory to fail to be axiological: (1) by making the right independent of considerations of goodness, or (2) by making the right depend on considerations of goodness, but not making it depend *solely* on such considerations.

Roughly speaking, a theory makes the right depend solely on considerations of goodness just in case there are certain features of choice situations the goodness of which completely determines which actions are permissible. For such a theory any two choice situations which have the same “goodness structure” (a vector of the goodness values associated with each of the specified features) also have the same “permissibility structure” (a vector of the permissibility values associated with each action) – no matter how different they may be with respect to their “non-value-theoretic” features. To make this notion clearer it will be helpful to explicitly identify and discuss a number of features that axiological theories may but need not have. The goal is not to be exhaustive, but rather to illustrate the great variety of ways in which the permissibility of actions may depend solely on considerations of goodness.

Axiological theories may, but need not, be *comparative* theories, i.e., theories for which the permissibility of a given action depends not only on its features, but also on features of the actions that are alternatives to it.¹¹ Act utilitarianism is an example of a comparative axiological theory. According to it the permissibility of a given action depends not only on the goodness of its consequences, but also on the goodness of the consequences of its alternatives. An action is permissible only if its consequences are at least as good as those of each of its alternatives. Minimal act utilitarianism – according to which an action is permissible just in case its

consequences are not bad – is an example of a non-comparative axiological theory. According to it the permissibility of a given action depends only on the goodness of its consequences. It does not depend on the goodness of the consequences – or any other features – of the other feasible actions.

Furthermore, even if comparative, axiological theories need not be *maximizing*, act utilitarianism is a maximizing axiological theory in that it judges an action permissible just in case its consequences are maximally good (at least as good as those of each of its alternatives). A variant of utilitarianism according to which an action is permissible just in case its consequences are at least as good as those of, say, *half* of its alternatives is an example of a non-maximizing comparative axiological theory.

Axiological theories may base the permissibility of actions on goodness from different *viewpoints*. Consider, for example, theories of the form: an action is permissible just in case it has the best consequences from the viewpoint of X. All such theories are axiological. X may be the agent (in which case we have ethical egoism), the agent's family, the agent's community, the universe as a whole (as, e.g., in classical utilitarianism), God, or whatever.

Axiological theories may even involve more than one viewpoint. Samuel Scheffler's¹² "hybrid" theory is an example of such an axiological theory. According to this theory an action is permissible just in case its consequences are from the *impartial* viewpoint at least as good as the best achievable state of affairs that does not require more than a specified level of sacrifice of goodness from the *agent's* viewpoint. Exactly how the level of sacrifice on the part of the agent is determined, and what the specified level is, need not concern us here. (Scheffler, in any case, is rather vague on this point.) The important point is that according to this theory the permissibility of actions depends only on the goodness of their consequences. It judges an action permissible just in case the goodness of its consequences from the impartial viewpoint is above a specified level, which is determined on the basis of the goodness from the agent's viewpoint of the consequences of the various feasible actions.

Axiological theories may, but need not, be *aggregative*, i.e., such that in determining the permissibility of actions the goodness from the viewpoint of some individual is traded off against that from the viewpoint of others. Total act utilitarianism is an aggregative axiological theory, since it directs the agent to maximize the *sum* of the goodness values of those affected. A theory that judges an action permissible just in case its consequences are *not bad* from the viewpoint of *any* person is a non-aggregative axiological theory. For such a theory there is no trading off of one person's good with that of another.

For axiological theories the permissibility of actions may depend not only on the goodness associated with the actions (be it that of the actions themselves, or that of associated state of affairs, such as their outcomes), but also on the goodness associated with other features of the choice situation. For example, a theory that judges an action permissible just in case its consequences are at least as good as *the status quo* is axiological. For this theory the permissibility of a given action

depends not only on the goodness of its consequences, but also on that of the status quo. Another example of such an axiological theory is a bargaining theoretic form of contractarianism according to which an action is permissible just in case its consequences are from the viewpoint of each person at least as good as their payoff under a specified bargaining solution (e.g., the Nash solution) given the feasible actions and the non-agreement outcome (e.g., the state of nature outcome, the status quo, or whatever).¹³

In summary, axiological theories are theories for which the permissibility of actions depends solely on considerations of goodness. They need not be maximizing, or even comparative. The goodness may be evaluated from person-relative viewpoints, an impartial viewpoint, or both. The goodness in question may be that of the action itself, its outcome, its anticipated outcome, its intended outcome, its reasonably anticipatable outcome, or whatever.

Just as there is a distinction between act and rule utilitarianism, there is a distinction between direct and indirect axiological theories. Direct axiological theories (such as act utilitarianism) ground the permissibility of actions in a given choice situation in and only in the goodness of features of *that choice situation*, whereas indirect axiological theories (such as rule utilitarianism) ground it in the goodness of features of other choice situations as well. For simplicity I have focussed, and shall continue to focus, on directly axiological theories.

I shall argue that the axiological/deontological distinction is more important than the teleological/non-teleological distinction as it is usually drawn. To do this it will be necessary to first clarify the exact nature of teleological theories.

4. Teleological theories

Teleological theories are a particular kind of axiological theory. As we shall see, however, authors disagree about the exact defining characteristics. Indeed, we shall uncover four different characterizations of teleological theories.

Almost all (if not all) authors require that a theory *maximize* the good in order to be teleological. As we shall see, most impose further requirements, but some do not. Nozick,¹⁴ and at least in certain passages, Rawls,¹⁵ characterizes teleological theories as theories that direct the agent to maximize the good, where the goodness in question may be that of the action's outcome, of its anticipated outcome and the like, or even of the action itself – no matter how this is evaluated (e.g., even if it is in terms of the agent's motives). This, then, is one way of characterizing teleological theories. Let us call theories that direct the agent to maximize the good *goodness maximizing*.

Most authors are more specific in their characterization of teleological theories. They require not merely that the theory direct the agent to maximize the good, but rather that they direct the agent to maximize the goodness of their outcomes, or of their intended, anticipated, or reasonably anticipatable outcomes.¹⁶ For brevity, let us use the term 'quasi-outcome' as a generic name for outcomes, in-

tended outcomes, anticipated outcomes, reasonably anticipatable outcomes, and the like. A second way of characterizing teleological theories, then, is as theories for which there is a specified kind of quasi-outcome such that an action is judged permissible just in case its specified quasi-outcome is maximally good. Let us call such theories *quasi-outcome-teleological*.

Many authors¹⁷ are even more specific in that they require the theory direct the agent to maximize the goodness of (objectively determined) outcomes – as opposed to the intended, anticipated, or reasonably anticipatable outcomes. Let us call such theories *outcome-teleological*. This, then, is a third way of characterizing theories.

There is yet another condition that is sometimes, but not always, invoked in the characterization of teleological theories: that the viewpoint from which the goodness of the outcomes is assessed be agent-invariant, i.e., the same no matter who the agent is. Classical act utilitarianism (which directs the agent to maximize the goodness of outcomes from the viewpoint of the universe as a whole) satisfies this condition, but ethical egoism (which directs the agent to maximize the goodness of outcomes from his/her viewpoint) and community utilitarianism (which directs the agent to maximize the goodness of outcomes from the viewpoint of his/her community) do not.

The distinction between teleological theories for which the viewpoint of evaluation is agent-invariant and those for which it is not is a special case of a more general (and more important) distinction between theories that are agent-sensitive and those that are not. A theory (teleological or not) is *agent-sensitive* just in case its permissibility conditions make an essential reference to the agent qua agent. Agent-sensitive theories may make the permissibility of actions depend in a special way on how they affect the agent and/or those bearing special relationships to the agent (e.g., friends, family, neighbors, etc.).¹⁸ Teleological theories are agent-insensitive just in case their viewpoint for the evaluation of outcomes is agent-invariant. As we saw above, some teleological theories are agent-sensitive (e.g., ethical egoism), and some are not (e.g., classical utilitarianism). Likewise, some non-teleological theories are agent-sensitive and some are not. For example, a theory that forbids the agent to kill or allow to die any member of *his/her* family is agent-sensitive, whereas a theory whose sole injunction forbids killing or allowing *anyone* to die is not.

Traditionally, agent-insensitivity has not been a defining characteristic of teleological theories. Ethical egoism (which is agent-sensitive) has generally been considered to be teleological.¹⁹ Some authors, such as Thomas Nagel and Samuel Scheffler,²⁰ do, however, impose the requirement of agent-sensitivity. A fourth way, then, of characterizing teleological theories is as theories that judge an action permissible just in case its outcome is maximally good from an agent-invariant viewpoint. Let us call such theories *agent-insensitive teleological*.

To sum up: We have seen that teleological theories have been characterized in the following four ways: (1) as goodness maximizing theories, i.e., as theories that judge an action permissible just in case it *maximizes* the good (no matter how this

is determined); (2) as quasi-outcome-teleological theories, i.e., as theories for which there is a specified type of quasi-outcome (outcome, intended outcome, anticipated outcome, reasonably anticipatable outcome, etc.) such that the theory judges an action permissible just in case its specified quasi-outcome is maximally good; (3) as outcome-teleological theories, i.e., as theories that judge an action permissible just in case its *outcome* (as opposed, e.g., to its anticipated outcome) is maximally good; (4) as agent-insensitive teleological theories, i.e., as theories that judge an action permissible just in case its outcome is *from an agent-invariant viewpoint* maximally good.

This raises the following question: Which, if any, of the above classes of theories provides the basis for the fundamental classificatory distinction? It is to this question that I now turn.

5. Fundamental classificatory distinctions

The teleological/deontological distinction is generally assumed to be the fundamental classificatory distinction for moral philosophy. Under the 'teleological/deontological distinction' title we have uncovered the axiological/deontological distinction and four teleological/non-teleological distinctions. I shall now argue that none of these distinctions is *the* fundamental classificatory distinction, because the presupposition that there is a unique fundamental distinction is false. There are rather several fundamental distinctions. Furthermore, I shall argue that none of the four teleological/non-teleological distinctions is even fundamental. The axiological/deontological distinction is significantly more important than any of them.

To say that a distinction is fundamental is to say that it is not significantly less important than any other distinction. Fundamental distinctions are ones that most ultimately matter. The importance of a distinction, I assume, is relative to a set of interests and purposes. Here we are concerned with the importance of distinctions relative to our interest in the assessment of moral theories. The importance of a distinction relative to these interests is determined by something like the usefulness of the role it plays, or would play on reflection, in the criticism and justification of moral theories.

The first question, then, is that of whether there is a unique fundamental classificatory distinction, i.e. a distinction that is significantly more important than any other. In order to answer this question, we need to determine whether there is a feature, the presence or absence of which is, relative to our interests in the assessment of moral theories, significantly more important than the presence or absence of any other feature.

A survey of some of the criticisms made of teleological theories supports fairly conclusively the thesis that there is no unique fundamental classificatory distinction. There are rather several distinct features that ultimately matter for theory assessment. For example, Ross²¹ criticizes certain kinds of teleological theories

for making the right depend solely on considerations of goodness. Moore²² replies at length to the criticism of his utilitarian theory that it bases the permissibility of actions on their outcomes, as opposed to the agent's motives (intended outcome), or the reasonably foreseeable outcome. Williams²³ criticizes certain kinds of teleological theories for making the permissibility of actions depend not only on what the agent "brings about", but also on what the agent allows to happen. Nagel²⁴ criticizes certain kinds of teleological theories for being agent-insensitive. Scheffler²⁵ criticizes certain kinds of teleological theories for being too demanding, in that they leave the agent too little liberty. Nozick²⁶ criticizes certain kinds of teleological theories for failing to respect the separateness of persons, in that they allow the interests of one person to be traded off against those of others. Mabbott²⁷ and other retributivists criticize certain kinds of teleological theories for being insensitive to the past.

Each of these authors identifies a specific feature and then criticizes certain kinds of teleological theories for having (or lacking) that feature. Each holds that such teleological theories are inadequate *because* they have (or lack) the feature in question. Indeed, each author holds that *any* theory – be it teleological or not – that has (or lacks) the feature in question is inadequate. Thus, each of these features provides the basis for a classificatory distinction that is important for the assessment of moral theories. Surely, at least two (if not most, or even all) of these distinctions are fundamental, i.e. not significantly less important than any other distinction. There is, that is, no unique fundamental classificatory distinction, but rather several of them.

As a minimum, then, the teleological/deontological distinction – however interpreted – has been overemphasized. Even if it is a fundamental distinction, it is not the only one.

But are any of the teleological/deontological distinctions that we have identified fundamental? The above survey of some of the features that ultimately matter for the justification and criticism of moral theories strongly supports a negative answer with respect to the four teleological/non-teleological distinctions. Moral philosophers do not criticize theories simply for being (or failing to be) teleological (in any of the four senses). They criticize them for having (or failing to have) more basic features, such as those identified above.

Consider, for example, the distinction between theories that are agent-insensitive teleological and those that are not. This is not a fundamental distinction, because being agent-sensitive teleological is not a feature that is invoked at the most basic level of criticism and justification of moral theories. Each of the following distinctions, for example, is significantly more important: (1) that between theories that are agent-sensitive and those that are not; (2) that between theories that make the right depend solely on considerations of goodness and those that do not; and (3) that between theories that make the right depend solely on their outcomes (and not, e.g., on their anticipated outcomes) and those that do not.

For the four teleological/non-teleological distinctions, the broader the characterization of teleological theories, the more it focuses on a feature of theories

that matters in an ultimate way, and the more important the corresponding teleological/non-teleological distinction is. The distinction between theories that are goodness maximizing and those that are not is the most important of the four distinctions. Even it, however, is not a fundamental distinction. Both (1) the distinction between theories that make the right depend solely on considerations of goodness and those that do not, i.e., the axiological/deontological distinction, and (2) the distinction between comparative and non-comparative theories are significantly more important.

The axiological/deontological distinction, on the other hand, is quite plausibly a fundamental distinction. As previously argued, the feature of making the right depend solely on considerations of goodness is a feature that figures in a basic way in the justification and criticism of moral theories.

6. Conclusion

In almost every introductory ethics course the student is taught that there are two basic kinds of moral theories: teleological and deontological. If the conclusions of this paper are correct, this practice must stop. As a minimum, discussion of “the” teleological/deontological distinction should be replaced by discussions of the axiological/deontological distinction. In addition, the student should be introduced to various of the other fundamental classificatory distinctions. This, I believe, will result in more penetrating theory criticism, and more creative theory construction. The hegemony of “the” teleological/deontological distinction must end.²⁸

NOTES

1. C.D. Broad, *Five Types of Ethical Theory* (London: Routledge & Kegan Paul, 1930), pp. 206, 278.
2. John Rawls, *A Theory of Justice* (Cambridge, Mass.: Belknap Press, 1971), p. 30.
3. See, for example, J.L. Mackie, *Ethics: Inventing Right and Wrong* (New York: Penguin Books, 1977), p. 159.
4. For example, C.D. Broad, *Five Types of Ethical Theory*, p. 206.
5. See, for example, William Frankena, *Ethics*, 2nd edition (Englewood, Cliffs: Prentice-Hall, 1973), p. 15.
6. See, for example, C.D. Broad, *Five Types of Ethical Theory*, p. 212.
7. See, for example, Alan Donagan, *The Theory of Morality* (Chicago: University of Chicago Press, 1977), Sec. 2.2, and Bernard Williams, “A Critique of Utilitarianism,” Sec. 3, in J.J.C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge, 1973).
8. See, for example, Charles Fried, *Right and Wrong* (Cambridge: Harvard 1978), p. 9, and C.D. Broad, *Five Types of Ethical Theory*, p. 278.
9. A characterization of this sort is given both by David Lyons, *The Forms and Limits of Utilitarianism* (London: Oxford 1965), p. vi, and by William Frankena *Ethics*, p. 15.
10. Michael Stocker, “Rightness and Goodness: Is There A Difference? *American Philosophical Quarterly* 10 (1973), 87–98. Stocker, however, is, for reasons that I cannot address here, skeptical of the significance of the axiological/non-axiological distinction.

11. There is, of course, the serious problem of explicating the difference between relational and non-relational features. Here I simply assume that there is a difference.
12. Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon Press, 1982), pp. 20–21.
13. David Gauthier, *Morals by Agreement* (London: Oxford University Press, 1986) advocates a theory roughly of this form.
14. Robert Nozick, *Philosophical Explanations* (Cambridge: Belknap Press of Harvard University Press, 1981), pp. 497–498.
15. For example, Rawls, *A Theory of Justice*, p. 24.
16. For example, Frankena, *Ethics*, p. 15.
17. For example, Donagan, *The Theory of Morality*, p. 51. Note that there are different types of outcomes that might be used as the basis for teleological theories. The *world scenario* of an action is the most complete (past, present, and future) state of affairs that would be realized if the action were performed. The *future* of an action is the most complete future state of affairs that would be realized if the action were performed. The *consequence* of an action is the most complete entirely avoidable state of affairs that would be realized if the action were performed. (A state of affairs is unavoidable for an agent in a given choice situation just in case it would be realized no matter what the agent does.) Some authors (such as Donagan) characterize teleological (consequentialist) theories as those that direct the agent to maximize the goodness of consequences in the above technical sense, whereas others allow the goodness to be of any type of outcome. In what follows I ignore for simplicity these further differences in the characterization of teleological theories. I discuss the significance of these differences in “Teleology, Consequentialism, and the Past”, forthcoming.
18. See Thomas Nagel, “The Limits of Objectivity,” in S. McMurrin, ed., *Tanner Lectures on Human Values* (Cambridge: Cambridge University Press, 1979) for further discussion of agent-sensitivity.
19. See, for example, Frankena, *Ethics*, p. 15–16.
20. See, for example, Nagel, “The Limits of Objectivity,” p. 119, and Scheffler, *The Rejection of Consequentialism*, p. 1. It should be noted that both of these authors use the term ‘consequentialist’ rather than ‘teleological’. Because these two terms are usually treated as synonymous, it is at least instructive to read them as yet another characterization of teleological theories.
21. W.D. Ross, *The Right and the Good* (New York: Bobbs-Merrill, 1930), Ch. 2.
22. G.E. Moore, *Ethics* (New York: Oxford University Press, 1912), Ch. 5.
23. Williams, “A Critique of Utilitarianism,” Sec. 3.
24. Nagel, “The Limits of Objectivity,” Sec. 3.
25. Scheffler, *The Rejection of Consequentialism*, for example, Ch. 2.
26. Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), Ch. 3.
27. J.D. Mabbott, “Punishment,” *Mind* 48 (1939), 152–167. Page references are to the reprint in Samuel Gorovitz, ed., *Mill: Utilitarianism* (Indianapolis: Bobbs-Merrill, 1971).
28. This paper is a condensed version of Chapter 2 of my dissertation “The Teleological/Deontological Distinction” (University of Pittsburgh, 1984). I am indebted to Kurt Baier, David Gauthier, Don Hubin, Shelly Kagan, Geoff Sayre McCord, and Nicolas Rescher for critical comments made on an earlier version of this paper.