

GIMMICKY REPRESENTATIONS OF MORAL THEORIES

PETER VALLENTYNE

1. Introduction

The teleological/deontological distinction is generally considered to be the fundamental classificatory distinction for ethics. I have argued elsewhere (Vallentyne 1987, and Ch. 2 of Vallentyne 1984) that the distinction is ill understood and not as important as is generally supposed. Some authors have advocated a more radical thesis. Oldenquist (1966) and Piper (1982) have both argued that the purported distinction is a pseudo-distinction in that any theory can be represented both as teleological and as deontological. Smart (1973, p. 13, and 1982) has also expressed views along these lines. Elsewhere (Vallentyne 1984, Ch. 3) I have shown that these arguments fail because the authors draw inadequate characterizations of the teleological/deontological distinction. Here I want to consider a challenge to the logical status of the distinction that raises deep and important questions about the structure of moral theories in general.

The challenge comes from an observation of Nozick's (1968, and 1974, pp. 28-39) concerning the distinction between maximizing theories, i.e., theories for which there is some complete, reflexive, transitive ranking relation¹ such that an action is judged permissible just in case it maximizes the ranking relation, and non-maximizing theories. His observation is that almost any theory can be given a maximizing representation (definition), and that it therefore seems that there is no basis for the intuitive distinction that we want to draw. Take any theory that intuitively seems to be non-maximizing, say the divine command theory (i.e., the theory that judges an action permissible just in case it violates none of God's commands). It is straightforward to give this theory a maximizing representation. Define a function, f , taking actions as arguments, such that $f(ac) = 1$ just in case ac satisfies God's commands, and $f(ac) = 0$ otherwise. The divine command theory is representable as a maximizing theory, since it judges an action permissible just in case it has a maximal f value.

So almost any theory, it seems, can be given a maximizing representation. Does this mean that the intuitive distinction between maximizing and non-maximizing theories is a pseudo-distinction?

¹ A relation, R , (e.g. "is at least as goods as") is *complete* just in case for any two objects, x_1 and x_2 , in its domain either: $R(x_1, x_2)$ or $R(x_2, x_1)$. A relation, R , is *reflexive* just in case for any object, x_1 , in its domain, $R(x_1, x_1)$. A relation, R , is *transitive* just in case: if $R(x_1, x_2)$ and $R(x_2, x_3)$, then $R(x_1, x_3)$.

Nozick does not think so. He believes that we need to distinguish between "gimmicky" and "natural" representations of theories. Maximizing theories are theories whose natural representation is maximizing – not simply those theories that *can* be given a (gimmicky) maximizing representation. Unfortunately, Nozick does not give us an account of the difference between gimmicky and natural representations. The question is whether there is a difference. If there isn't, then the intuitive maximizing/non-maximizing distinction is illusory.

Similarly, it would seem (although we shall have to return to this point) that almost any theory can be given a teleological representation. Consequently, it would seem that unless there is a difference between gimmicky and natural representations of theories, the intuitive teleological/deontological distinction is also illusory.

2. The Problem Analyzed

Before considering this problem in greater detail, a few terminological conventions need to be made explicit.

First of all, moral theories may make pronouncements on a variety of moral matters: for example, the permissibility of actions, the permissibility of social institutions, the goodness of states of affairs, the goodness of actions, the goodness of character traits, etc. The teleological/deontological distinction concerns theories of *permissibility*. For simplicity we are focussing on the distinction as it applies to theories of permissibility for actions – as opposed to, e.g., social institutions. Therefore, except where otherwise noted, "moral theory" is to be understood as short for "moral theory of the permissibility of actions".

Second, moral theories connect moral concepts to other concepts. For example, act utilitarianism connects the concept of permissibility for actions to the concept of having a maximally good outcome, and the divine command theory connects the concept of permissibility for actions to the concept of obeying God's commands. There are different metaethical positions concerning the nature of this connecting relations. Some hold that it is some sort of analytic relation, such as conceptual equivalence (e.g., Moore 1903, pp. 146–48), explication (e.g., Rawls 1971, p. 111), or reforming definition (e.g., Brandt 1979, ch. 1). Others (e.g., Ross 1930, ch. 1, and Ross 1939, pp. 26–28) hold that it is some sort of synthetic relation, such as that which provides the reasons or grounds for a moral concept holding when it does (that which *makes* actions permissible). On the former view moral theories are about the *essence* of moral concepts, whereas on the latter view they are about their *grounds*.

Fortunately, for the present purposes we need not determine the exact nature of this connection. The important point is that moral theories (of permissibility) posit some sort of connection between the

concept of permissibility and some other concept. Let us call the concept that a theory connects to the concept of permissibility its *basis* for the concept of permissibility for actions.

A *representation* of a theory is a concept that has the same satisfaction conditions as the theory's basis. The representing concept and the theory's basis are satisfied by exactly the same actions in all possible choice situations. The problem raised by Nozick is that almost any theory can be given a maximizing representation, i.e., for almost every theory a concept can be found that has the same satisfaction conditions as the theory's basis and that is definable as maximizing some complete, reflexive, transitive ranking relation. This suggests that, if our intuitive distinction between maximizing and non-maximizing theories is sound, there must be a difference between natural representations and gimmicky representations. The question is whether there is a difference.

Nozick's remarks suggest that *any* theory can be given a maximizing representation. Strictly speaking, this is not true – at least not if it is assumed that there are only a finite number of feasible actions. Under this assumption, for any maximizing theory there is always at least one action that is ranked maximally good (at least as good as all the others), and therefore judged permissible. Consequently, theories that allow the possibility of *prohibition dilemmas*², i.e., choice situations in which *no* action is permissible, cannot be given a maximizing representation. (A theory that forbids one to kill or allow to die any member of one's family would give rise to a prohibition dilemma in a situation where unless you kill your father he will kill your mother.) So some theories cannot be given a maximizing representation.

There is yet a further reason to question the claim that *any* theory can be given a maximizing representation – at least if one assumes both that no action is part of more than one choice situation, and that the ranking relation to be maximized must be *choice situation invariant* (e.g., the same for all agents and all times). Consider, for example, a theory that: judges a1 but not a2, permissible in one choice situation; judges a2, but not a3 permissible in a second choice situation; and judges a3, but not a1, permissible in a third situation.³ This theory cannot be given a maximizing representation, because: by the first choice situation a1

² In Valentyne forthcoming, I distinguish prohibition dilemmas (choice situations in which every action is forbidden) from obligation dilemmas (choice situations in which more than one action is obligatory), and argue that the former, but not the latter, are conceptually possible.

³ For example, consider a theory that (1) directs the agent to save the life of the *older* boy, when confronted with the choice of saving the life of one of two boys, and (2) directs the agent to save the life *healthier* child, when confronted with the choice of saving either a boy or a girl. Let the first choice situation be the choice between saving Unhealthy Older Boy (a1) and Very Healthy Younger Boy (a2); let the second choice situation be the choice between saving Very Healthy Younger Boy (a2) and Moderately Healthy Girl (a3); and let the third choice situation be the choice between saving Moderately Healthy Girl (a3) and Unhealthy Older Boy.

would have to be ranked higher than a_2 ; by the second choice situation a_2 would have to be ranked higher than a_3 ; and by the third choice situation a_3 would have to be ranked higher than a_1 . Because such a ranking relation would not be transitive (since $a_1 > a_2 > a_3$, and $a_3 > a_1$), this theory cannot be given a maximizing representation.

The procedure specified above (involving the gimmicky f function) for producing a maximizing representation implicitly assumed that actions are individuated in such a way that no action is feasible in more than one conceptually possible choice situation. This assumption (or alternatively, the assumption the ranking relation need not be choice situation invariant) greatly facilitates the representation of theories as maximizing theories. On this assumption all that matters is that *in a given choice situation* all and only the permissible actions are ranked maximally good. The ranking of actions in different choice situations does not matter, since by assumption they are distinct actions, and so can be arbitrarily ranked any way whatsoever.

Of course, if actions are individuated in such a way that they cannot be part of more than one conceptually possible choice situation, or if the ranking relation to be maximized need not be choice situation invariant, then this second way of showing that not all theories can be given a maximizing representation does not get off the ground. In fact, because I accept both of the antecedents, I find this second line of attack ineffective. Rather than argue the point, however, let us grant it for the sake of argument. For, even if, in addition to the first line of attack (concerning theories that allow prohibition dilemmas), this second line of attack is successful, it still remains true that many intuitively non-maximizing theories can be given a maximizing representation. Thus, we cannot avoid the issue of what, if anything, the difference is between natural and gimmicky representations of theories.

It might be suggested that the natural representation of a theory is provided by the *definition* of its basis. The natural representation of a theory is a maximizing representation (and the theory is truly a maximizing theory), for example, if and only if the theory's basis has a maximizing definition (it being defined as maximizing some ranking relation). The problem with this suggestion is that concepts, at least as they are often thought of, do not have privileged definitions. Any given concept can be defined in a variety of ways. As we saw above the concept of satisfying God's commands can be given a maximizing definition in terms of the function f . Or consider the concept of bachelorhood. We usually think of this concept as having the privileged definition of being the conjunction of unmarriedness with manhood. But it can be defined in many other ways. It can, for example, be defined in terms of dentishood. A simple truth-table check shows that bachelorhood is conceptually equivalent to the concept (dentist-or-dachelor)-and-bentist where a dachelor is something that is a bachelor

but not a dentist, and a bentist is something that is either a bachelor or not a dentist.

Furthermore, not only can every concept be defined in a variety of different ways, every concept can be defined in terms of *any arbitrarily chosen* concept. One just has to cook up the right defining clause (as in the above definition of bachelorhood in terms of dentishood). The question of whether a given concept can be defined in terms of another in conjunction with the members of some *relatively restricted* set of concepts can be any interesting question, but if no restriction is placed on the concepts that may be used in the definition, the question will always receive an affirmative answer.

Concepts have satisfaction conditions that divide the objects of the world into those that do and those that do not satisfy them. It seems, however, that they do not have any structure. And so it doesn't seem possible to distinguish between "natural" and "gimmicky" definitions of concepts.

3. A Possible Solution

The problem of distinguishing between natural and gimmicky representations of moral theories is simply a special case of the more general problem of distinguishing between natural and gimmicky representations (definitions) of concepts. If concepts have privileged (natural) definitions, then moral theories have privileged (natural) representations, namely the privileged definition of the concept that is their basis. Recognizing this suggests a solution to the problem. In the preceding subsection I assumed that concepts can be defined in a variety of ways, no one of which is privileged. The time has come to question this assumption.

Recently authors have started distinguishing between two kinds of concepts (and propositions): coarse-grained and fine-grained. Coarse-grained concepts are individuated solely in terms of their satisfaction conditions. They do not have any structure. They do not have other concepts as constituents. They can be defined in a variety of ways. Everything said above about concepts is true of coarse-grained concepts. In particular, there is no relevant distinction between gimmicky and natural representations of coarse-grained concepts.

Fine-grained concepts, on the other hand, do have structure and constituents. The fine-grained concept expressed by "unmarried man" is not the same as the fine-grained concept expressed by "(dentist-or-dachelor)-and-bentist" – even though both express the same coarse-grained concept (in that both have the same satisfaction conditions). The former concept has the concepts of unmarriedness and of manhood as parts, whereas the latter does not. David Lewis (1972) has given a very elegant explication of the notion of fine-grained concepts in terms

of set-theoretical constructions of coarse-grained concepts. More specifically, he explicates them as finite ordered trees of coarse-grained concepts. Simplifying somewhat, this amounts to thinking of fine-grained concepts as certain types of ordered sets of coarse-grained concepts. Thus, for example, the fine-grained concept of being an unmarried man might be thought of as the ordered triple \langle conjunction, unmarried, man \rangle . The structure of the concept – as given by the first element – is conjunctive, and the concept has two constituent concepts: unmarriedness and manhood.

Fine-grained concepts are individuated not only in terms of their satisfaction conditions but also in terms of their structure and their constituents. A gimmicky representation of a fine-grained concept is one that captures its satisfaction conditions, but not its fine-grained structure.

Likewise, fine-grained propositions are individuated not only in terms of their truth conditions (as coarse-grained propositions are), but also in terms of their structure and their constituents. A gimmicky representation of a fine-grained proposition is one that captures its truth conditions, but not its fine-grained structure.

There is nothing strange or weird about fine-grained concepts and propositions. As David Lewis has shown, they can be explicated as certain types of set-theoretical constructions out of coarse-grained counterparts. Furthermore, there are precedents for making use of them to solve philosophical problems. David Lewis (1972) has used them to give an account of the meanings of linguistic items. The merit of this approach is that sentences with the same truth conditions, need not be interpreted as having the same meaning. For example, although "Grass is green or it isn't" and "Snow is white or it isn't" express the same coarse-grained proposition, they do not express the same fine-grained proposition. Using fine-grained concepts and propositions Lewis is able to capture the intuition that the two sentences have different meanings. And Hartry Field (1980) has discussed the use of fine-grained propositions as the objects of belief. The advantage of such an approach is that one can distinguish between someone believing that $2 + 2 = 4$, and he/she believing other theorems of arithmetic. Since all the theorems of arithmetic have the same truth-conditions (they are all necessarily true), they all express the same coarse-grained proposition. Using fine-grained propositions as the objects of belief enables one to capture the intuition that believing that $2 + 2 = 4$ is not the same as believing some deep and surprising mathematical theorem.

We can now see that the problem of distinguishing natural and gimmicky representations of moral theories arises if we think of the bases of moral theories (i.e., the concepts that they connect to the concept of permissibility) as being *coarse-grained*. Because there is no distinction that can be made between gimmicky and natural representa-

tions of coarse-grained concepts, if moral theories are so construed, there is no basis for the intuitive distinction between gimmicky and natural representations of moral theories.

If, on the other hand, we think of the bases of moral theories as fine-grained concepts, the distinction between gimmicky and natural representations of theories is straightforward. The natural representation of a theory is given by the fine-grained representation of its basis. Gimmicky representations are representations that have the same satisfaction conditions, but not the same fine-grained constitution as the natural representation. Thus, the coarse-grained concept designated by "satisfies God's commands" is the same as that designated by "has a maximal f value for the set of feasible actions" (where f is the function defined over actions that takes the value 1, if they satisfy God's commands, and 0 otherwise). The fine-grained concepts designated by these expressions are not, however, the same. The latter, but not the former, has a maximizing structure. Since the former fine-grained concept represents the basis of the divine command theory as we intuitively understand it, it is the natural representation of that theory. The latter fine-grained concept does indeed have the same satisfaction conditions, but its fine-grained structure does not represent the divine command theory. It is therefore merely a gimmicky representation.

Thus, if we construe the bases of moral theories as fine-grained concepts, the distinction between natural and gimmicky representations of moral theories can be made. Note, however, that although I have shown that this distinction *can* be made, I have not shown that it is *worth* making. It might be objected, after all, that the natural/gimmicky distinction is irrelevant for moral philosophy, that all that really matters in assessing moral theories are the satisfaction conditions of their bases. A full discussion of this issue is beyond the scope of this paper, so I shall only make a few brief remarks.

Whether or not the distinction is worth making depends presumably on whether there are fine-grained facts of the type represented by moral theories. For example, if moral theories are interpreted as representing an objective (mind-independent) moral order, then presumably the distinction is worth making if and only if there are objective fine-grained facts (in virtue of which one fine-grained theory would be true, and all others false). Because I am inclined to think that there are no such facts, I am inclined to think that the natural/gimmicky distinction is not worth making for theories so interpreted. If there is no way of distinguishing the truth of one fine-grained representation from that of other coarse-grained equivalents, there is no point in making the distinction.

On the other hand, if moral theories are interpreted as representing the considered moral judgments of particular persons (it being claimed that the theory judges an action permissible just in case the person does), then presumably the distinction is worth making if and only if

people distinguish different fine-grained representations of theories. Because people do make such distinctions (e.g., between a usual representation of the divine command theory and the maximizing representation suggested above), the natural/gimmicky distinction may be worth making for moral theories so interpreted. Just as we want to distinguish between someone's believing that $2 + 2 = 4$ from believing some deep and surprising truth of arithmetic, we would want to distinguish between different fine-grained representations of a person's coarse-grained moral judgments. Making the distinction would enable us, for example, to explain why a person would assent to the sentence "An action is permissible just in case it satisfies God's Commands", but not to "An action is permissible just in case it has a maximal f -value" (where f is the gimmicky function described above).

4. Implications for the Teleological/Deontological Distinction

So far we have focussed on the maximizing/non-maximizing distinction. Many intuitively non-maximizing theories can be given a maximizing representation. Consequently, unless we can distinguish between natural and gimmicky representations of moral theories, there is no basis for the intuitive maximizing/non-maximizing distinction. As we shall now see, this is not so for the teleological/deontological distinction.

Teleological theories are a particular kind of maximizing theory: they judge an action permissible just in case its outcome is maximally good. (I restrict my attention here to *act* teleological theories.) Unlike the case for maximizing theories, it is not sufficient for a theory to be teleological that it direct the agent to maximize *some* ranking relation. The ranking relation must be that of the *intrinsic goodness* for states of affairs. Furthermore, the goodness associated with an action must be that of its *outcome* (understood as including, perhaps, the performance of the action). The defining characteristics of teleological theories are significantly more demanding than those of maximizing theories. Consequently, it is not as easy for a theory to be given a gimmicky coarse-grained teleological representation. This suggests that, unlike the case for the intuitive maximizing/non-maximizing distinction, the intuitive teleological/deontological distinction may be captured by the associated coarse-grained distinction.

The following line of reasoning suggests that it does. The defining characteristic of teleological theories is that they direct the agent to maximize the goodness of outcomes. Many intuitively deontological theories, however, do not even have a theory of the good, and for those that do it is unlikely that they judge an action permissible just in case its outcome is maximally good. For example, the divine command theory coupled with almost any theory of the good cannot be given a teleological representation. Such a theory, can, it is true, be represented

as directing the agent to maximize *some* ranking relation, but, for almost any theory of the good that is joined to it, it cannot be represented as directing the agent to maximize the goodness of outcomes as judged by that theory of the good. Thus, because intuitively deontological theories cannot be given coarse-grained teleological representations and intuitively teleological theories can, the coarse-grained teleological/deontological distinction does – this line of reasoning suggests – capture the intuitive distinction.

In response to this line of reasoning the following objection can be raised. Although it is true that *most* intuitively deontological theories cannot be given coarse-grained teleological representations, *some* can be. The above reasoning points out that most intuitively deontological theories either do not have theories of the good, or, if they do, their theories are not such that an action is judged permissible just in case its outcome is maximally good. What this reasoning overlooks, it is now claimed, is that for many of these theories a theory of the good can be provided (replacing their old theory, if they had one) such that relative to the new theory of the good they can be given a coarse-grained teleological representation. The same types of tricks that can be used to give almost any theory a maximizing representation can be used to show that at least some intuitively deontological theories can be given coarse-grained teleological representations. One just has to go a step further and specify that the gimmicky ranking relation represents the theory's theory of the good. Because theories so produced still have intuitively deontological theories of the right, this shows, it is claimed, that the coarse-grained teleological/deontological distinction does not capture the intuitive distinction.

Both this objection and the initial line of reasoning are based on a misunderstanding of what it takes to be teleological. They both assume that a theory is teleological if and only if it judges an action permissible just in case the action's outcome maximizes *that which the theory holds to be the basis of goodness*. This, however, is a mistake. There is an important difference between (1) directing agents to maximize *some* ranking relation, that the theory *happens* to hold to be the basis of goodness, and (2) directing agents to maximize goodness, *however* that may be determined. A theory must be of the latter type in order to be teleological.

Being teleological is like directing one's son to marry the richest woman in town – *whoever* that may be. For any particular woman, one's judgments about whether he should marry her, depend on one's judgment about how rich she is. Being teleological is *not*, on the other hand, like directing one's son to marry Ms. Jones, whom one *happens* to believe to be the richest woman in town. One's judgments about whether he should marry any particular woman do not in that case depend on one's judgment about how rich she is.

Teleological theories judge an action permissible just in case it maximizes the good – whatever the exact basis of goodness is. Their judgments of permissibility depend on their theory of the good. In order to determine the permissibility of an action, they need to use their theory of the good. Change their theory of the good, and you change their judgments of permissibility.

Thus, the fact that some intuitively deontological theories can be represented as directing the agent to maximize *what the theory takes to be the basis of goodness* in no way shows that the coarse-grained teleological/deontological distinction fails to capture the intuitive distinction. Such representations are not coarse-grained teleological, because they do not represent the theory as judging actions permissible just in case *goodness* – whatever its exact basis – is maximized. Rather, they represent the theory as judging actions permissible just in case the *basis of goodness* – according to a *particular* theory of the good – is maximized. The gimmicky representations do not make the theories' judgements of permissibility depend on their judgements of goodness. Changing the theories of the good associated with each theory of permissibility does not change the judgements of permissibility of particular actions. Thus, because they do not have the appropriate satisfaction conditions, the purportedly coarse-grained teleological representations of intuitively deontological theories are thus not truly coarse-grained teleological.

In summary, Nozick's observation that almost every theory can be given a maximizing representation does not threaten the intuitive teleological/deontological distinction. Unlike the case for the maximizing/non-maximizing distinction, the teleological/deontological distinction does not rely on a fine-grained interpretation of moral theories. To be coarse-grained teleological a theory must direct the agent to maximize the goodness of outcomes however that may be determined. That is a demanding requirement that only intuitively teleological theories satisfy.⁴

University of Western Ontario
London, Ontario
Canada N6A 3K7

⁴ This paper is drawn from Ch. 3 of Vallentyne 1984. For helpful comments on an earlier version of this paper I would like to thank David Gauthier, Shelly Kagan, Al Roth, and Richard Sikora.

References

Brandt, Richard B. (1979). *The Theory of the Good and the Right* (Oxford: Oxford University Press).

- Field, Hartry (1980) "Mental Representations", in Ned Block, ed. *Readings in Philosophy of Psychology* (Cambridge, Mass.: Harvard University Press, 1980).
- Lewis, David (1972), "General Semantics", in Davidson and Harman, eds. *Semantics of Natural Languages* (Dordrecht-Holland: D. Reidel, 1972).
- Moore, G. E. (1903). *Principia Ethica* (London: Cambridge University Press).
- Nozick, Robert (1968). "Moral Complications and Moral Structures", *Natural Law Forum*, 13:1–50.
- Nozick, Robert (1974). *Anarchy, State, and Utopia* (New York: Basic Books).
- Oldenquist, Andrew (1966). "Rules and Consequences", *Mind*, 75: 180–192.
- Piper, Adrian (1982). "A Distinction Without a Difference", in French, Uehling, and Wettstein, eds. *Midwest Studies in Philosophy, Vol. 7, Social and Political Philosophy* (Minneapolis: University of Minnesota Press, 1982), pp. 403–435.
- Rawls, John (1971). *A Theory of Justice* (Cambridge, Mass: Belknap Press).
- Ross W. D. (1930). *The Right and the Good* (New York: Oxford University Press)
- Ross, W. D. (1939). *Foundations of Ethics* (London: Oxford University Press).
- Smart, J. J. C. (1973). "An Outline of a System of Utilitarian Ethics", in Smart and Williams, eds. *Utilitarianism: For and Against* (Cambridge: Cambridge U.P., 1973), pp. 3–74
- Smart, J. J. C. (1982). "Hedonistic and Ideal Utilitarianism", in French, Uehling, and Wettstein, eds. *Midwest Studies in Philosophy, Vol. 7* (Minneapolis: University of Minnesota Press, 1982), pp. 240–251.
- Vallentyne, Peter (1984). "The Teleological/Deontological Distinction", unpublished dissertation, University of Pittsburgh.
- Vallentyne, Peter (1987). "The Teleological/Deontological Distinction", *Journal of Value Inquiry*, 21: 21–32.
- Vallentyne, Peter, (forthcoming). "Two Types of Moral Dilemmas" *Erkenntnis*.