

## Reflections on Consequentialism

by

Lars Bergström

(Stockholm University)

Consequentialism, or utilitarianism (in a broad sense), can be interpreted in many different ways.<sup>1</sup> In *The Cambridge Dictionary of Philosophy*, there is no separate entry for "consequentialism"; the reader is referred to "utilitarianism", and this is in turn explained as "the moral theory that an action is morally right if and only if it produces at least as much good (utility) for all people affected by the action as any alternative action the person could do instead".<sup>2</sup> This is a sufficient starting point for my purposes in this paper. It indicates, for example, that the kind of theory I have in mind is "act consequentialism", rather than "rule consequentialism".

A well-known consequentialist theory would be the hedonistic utilitarianism described in the first two chapters of G.E. Moore's *Ethics*; another example would be Moore's own non-hedonistic "ideal" utilitarianism proposed in the same book.<sup>3</sup> Moore takes (hedonistic) utilitarianism to contain the "principle" that "a voluntary action is right, whenever and only when the agent could *not*, even if he had chosen, have done any other action instead, which would have caused more pleasure than the one he did do".<sup>4</sup> Notice, however, that, for Moore, utilitarianism does not only say "that the producing of a maximum of pleasure is a characteristic, which did and will belong, *as a matter of fact*, to all right voluntary actions (actual or possible)", it also says that "it is *because* they possess this characteristic that such actions are right".<sup>5</sup> I follow Moore here. I take consequentialism to be an explanatory theory of moral rightness; it is neither a statement of a mere correlation between rightness and a certain empirical and/or evaluative property, nor a decision procedure to be followed in actual situations of choice.

The distinction between theories of rightness and practical decision procedures is probably quite old. It appears to go back at least to Henry Sidgwick.<sup>6</sup> It is quite clearly

---

<sup>1</sup>See e.g. Dan W. Brock, "Recent Work in Utilitarianism", *American Philosophical Quarterly*, 10: 4 (October 1973): 241-76; Lars Bergström, "On the Formulation and Application of Utilitarianism", *Nous*, 10 (1976): 121-44; and Erik Carlson, *Consequentialism Reconsidered*, Theory and Decision Library 20 (Dordrecht: Kluwer Academic Publishers, 1995).

<sup>2</sup>*The Cambridge Dictionary of Philosophy*, ed. by Robert Audi (Cambridge: Cambridge University Press, 1995), p. 824.

<sup>3</sup>G.E. Moore, *Ethics* (London: Oxford University Press, 1912).

<sup>4</sup>*Ibid.*, p. 19.

<sup>5</sup>*Ibid.*, p. 29.

<sup>6</sup>See *The Methods of Ethics*, 7th Edition (London: Macmillan, 1967), where Sidgwick distinguishes between "ultimate reasons" for acting and "methods" for determining right conduct (e.g., p. 78), and between "principles" and "methods" (e.g., pp. 119 and 121). See also the discussion in Samuel Scheffler, *The Rejection*

formulated, for example by Eugene Bales in 1971 and by Peter Railton in 1987. Bales stresses "the importance of maintaining a sharp distinction between (a) decision-making procedures, and (b) accounts of what makes right acts right".<sup>7</sup> Railton distinguishes, in a similar manner, between subjective and objective consequentialism. He writes:

*Subjective consequentialism* is the view that whenever one faces a choice of actions, one should attempt to determine which act of those available would most promote the good, and should then try to act accordingly. [...] *Objective consequentialism* is the view that the criterion of rightness of an act or course of action is whether it in fact would most promote the good of those acts available to the agent".<sup>8</sup>

Railton also points out that "objective consequentialism sets a definite and distinctive criterion of right action, and it becomes an empirical question (though not an easy one) which modes of decision making should be employed and when".<sup>9</sup>

It seems to be a fairly common idea that appropriate decision procedures, or modes of decision making, have to be determined on the basis of some theory or principle of right action. For all we know, it may be disastrous—and morally wrong—to try to apply the criterion of rightness itself as a rule for decision-making in ordinary life. If a given principle of rightness is easy to apply in ordinary life, it may also be appropriate as a decision procedure. But if a principle is difficult to apply, e.g. because its application requires factual information which is very difficult or even impossible to obtain within the time limits of ordinary decision making—and a consequentialist principle is clearly of this kind—then some more practical decision rule has to be used. It is often said that we need "rules of thumb" here. Moreover, different decision procedures may be appropriate for different people and in different types of situation. As Sidgwick points out, "it may be best on the whole that there should be conflicting codes of morality in a given society at a certain stage of its development".<sup>10</sup> In some situations, it might be best to use the "rule" that no explicit rule at all should be used, so that the agent can just do what he likes or what first occurs to him.

As far as I know, consequentialism is usually proposed as a theory of rightness, not as a decision procedure or method of deliberation. In Railton's terms, objective consequentialism is more common than subjective consequentialism. And objective consequentialism is more

---

*of Consequentialism* (Oxford: Clarendon Press, 1994), pp. 44-50, on the distinction between justification and motivation in consequentialist thinking.

<sup>7</sup>R. Eugene Bales, "Act-utilitarianism: account of right-making characteristics or decision-making procedure?" *American Philosophical Quarterly*, 8: 3 (July 1971): 257-65, p. 257.

<sup>8</sup>Peter Railton, "Alienation, Consequentialism, and the Demands of Morality", *Philosophy & Public Affairs*, 13: 2 (April 1984): 134-71, p. 152.

<sup>9</sup>*Ibid.*, p. 156.

<sup>10</sup>*The Methods of Ethics*, p. 491.

often discussed by moral philosophers than is subjective consequentialism. I think that this is a pity. It seems to me that some form of subjective consequentialism is more plausible than objective consequentialism. In order to see this, let us now survey the pros and cons of objective consequentialism.

### **1. Reasons in favour of objective consequentialism**

The most important question concerning objective consequentialism is perhaps: Is it *true*? It is not easy to answer this question. You cannot directly *see* whether it is true or false. Some people find it very plausible, others find it very implausible. You cannot directly see whether a given scientific theory is true either, but in this case there are certain more or less uncontroversial criteria for theory-choice, which scientists can and do apply in order to evaluate the given theory. Typically, such criteria state good-making or bad-making characteristics that scientific theories may have or lack. For example, Thomas Kuhn has pointed out that empirical adequacy, consistency, broad scope, simplicity, and fruitfulness are widely accepted as good-making characteristics.<sup>11</sup> It would be interesting to have a similar list of widely accepted good-making characteristics for moral theories of rightness, but I shall not now try to make up such a list. Instead, I shall mention five good-making characteristics which seem to apply to consequentialism, and which seem to constitute the main reasons in favour of this theory. I myself do not wish to say, without reservation, that all of these reasons are really valid, but I have some inclination to say this, and I think it is what objective consequentialist have said or might say in favour of their view. In the next section, I will mention five bad-making characteristics, which pull in the opposite direction.

(i) *Consequentialism is morally attractive*. In the first place, it may even seem intuitively self-evident that we should always act so as to achieve the best possible outcome. What could be better? Besides, if we think more specifically of classical utilitarianism, as a main version of consequentialism, this is morally attractive in at least two different respects. First, it stresses the way people (or sentient beings, in general) are affected by our actions. What matters is the welfare or the preferences of everyone to whom our actions make a difference. This is surely very important. It expresses a generalized version of the idea that one should treat others as one would like others to treat oneself. Second, it is completely impartial. It rejects all forms of egoism, ethnocentrism, and racism. It is also temporally impartial. Future generations have the same weight in a consequentialist calculus as our own generation (other things being equal).

---

<sup>11</sup>See e.g. T.S. Kuhn, "Objectivity, Value Judgment, and Theory Choice", in *The Essential Tension* (Chicago: University of Chicago Press, 1977), p. 322. It might be supposed that the characteristics mentioned by Kuhn should be interpreted as methodological rules, inductively related to the aims of science, rather than as good-making characteristics (constituting, as it were, the aims of science). I have argued, however, that such an interpretation would be mistaken; see Lars Bergström, "Scientific Value", *International Studies in the Philosophy of Science*, 10:3 (1996), 189-202.

(ii) *Consequentialism solves moral conflicts.* In ordinary life, moral considerations sometimes point in different directions. Those are the very situations in which a moral theory is most needed. When different moral rules give different directions, there is a moral conflict. We need to know how to handle such conflicts. Consequentialism presents a general solution to such conflicts. This solution may be hard to identify in practice, but it is at least a solution. According to consequentialism there is always an answer to hard moral questions.

(iii) *Consequentialism is a bold conjecture.* For one thing, it is bold in the sense that it is simple and has very broad scope. Maybe it can also be said to be "bold" in a sense similar to that stressed by Karl Popper for scientific theories,<sup>12</sup> namely that it contradicts earlier theories while at the same time explaining their relative success. Consequentialists often claim that many moral principles, which are strictly speaking incompatible with consequentialism, can actually be given a consequentialist motivation if they are interpreted as useful approximations to be used in practice. Also consequentialism can be said to be bold in the sense that it has more content than many alternative views. Alternative theories, such as Kantianism, Christian ethics, existentialism, natural right theories, and so on, seem to give rather indeterminate answers to actual moral problems. Consequentialism, on the other hand, has a definite answer to every question concerning the moral rightness of actions.

(iv) *Consequentialism is theoretically fruitful.* In the development of moral philosophy since the time of Sidgwick, say, utilitarianism and consequentialism in general has played a major part. It has been theoretically very fruitful in the sense that it has stimulated philosophers to work out details and answer difficult objections. It has given rise to many interesting problems and to a lot of professional discussion. In recent years, John Rawls's theory has also been very influential, but on the whole I think it is fair to say that consequentialism has been more theoretically fruitful, in our time, than any other moral theory.

(v) *Consequentialism is our best theory so far.* Even if consequentialism is not a completely satisfactory theory, it nevertheless seems to be better than the currently available alternative theories of rightness. It is better, I think, in virtue of the considerations (i) through (iv) above. Moreover, it seems reasonable to hold on to a theory which is better than any known alternative, even if the theory itself is problematic in many ways. This seems to be a good strategy in science, and I think it can be applied in ethics as well (provided, of course, that the theory solves some important problems). It is better to have a theory with some advantages than no theory at all, and the more advantages the better.

## **2. Reasons against objective consequentialism.**

---

<sup>12</sup>See e.g. Karl Popper, *Objective Knowledge* (Oxford: Clarendon Press, 1972), p. 16.

A great many objections have been raised against the usual versions of consequentialism and utilitarianism. Even the proponents of consequentialism would agree that some of these objections are serious, at least at first sight. Whether or not they are fatal is a moot point. The objections can be systematized as follows.

(i) *Consequentialism is morally offensive*. The most common type of objection to consequentialism consists in the construction of some thought experiment in which consequentialism appears to have morally unacceptable consequences. For example, consequentialism seems to imply that, under certain conditions, you ought to kill an innocent person in order to save several others,<sup>13</sup> people should be punished for crimes they have not committed, organs for transplantation should be taken from healthy people against their will,<sup>14</sup> and so on. It is unclear, however, what is shown by such objections. Utilitarians typically answer them by invoking the distinction between theories of rightness and decision procedures.<sup>15</sup> They hold that objections of this kind is no threat to *objective* consequentialism.

(ii) *Consequentialism is too demanding*. Since consequentialism requires that we always act so as to bring about the best possible outcome, we can be fairly sure that we *never* do what is right according to consequentialism. For example, instead of going to the cinema, I could have used the money to save the life of starving people in some poor country. Now, it may not be unreasonable to hold that most or all of my actions are morally wrong. But it is less reasonable to say that in all probability my actions would still be wrong, even if I led a much better life in consequentialist terms. However, this seems true, given consequentialism.<sup>16</sup> For according to consequentialism, *everything* except the very best is morally wrong. This seems absurd.

(iii) *Consequentialism is incomplete*. It seems that every plausible version of consequentialism would have to involve interpersonal comparisons of welfare levels and/or of welfare differences.<sup>17</sup> Unless such comparisons can be made in an objective way, at least in principle, consequentialism is incomplete in the sense that it does not give definite answers to all questions concerning the rightness of actions. Unfortunately, the dominant opinion among specialists seems to be that the problem of interpersonal comparisons has not yet been solved. I myself used to think that the problem could be solved,<sup>18</sup> but I must say in

---

<sup>13</sup>See Bernard Williams, "A Critique of Utilitarianism", in J.J.C. Smart and B. Williams, *Utilitarianism: for and against* (London: Cambridge University Press, 1973), p. 98f.

<sup>14</sup>See e.g. Gilbert Harman, *The Nature of Morality* (New York: Oxford University Press, 1977), p. 3f.

<sup>15</sup>See e.g. R.M. Hare, *Moral Thinking* (Oxford: Clarendon Press, 1981), p. 131f.

<sup>16</sup>Objections of this kind are discussed in Björn Eriksson, *Heavy Duty: On the Demands of Consequentialism* (Stockholm: Almqvist&Wiksell, 1994), where further references can be found.

<sup>17</sup>See e.g. Michael D. Resnik, *Choices: An Introduction to Decision Theory* (Minneapolis: University of Minnesota Press, 1987), pp. 205-12.

<sup>18</sup>See Lars Bergström, "Interpersonal Utility Comparisons", *Grazer Philosophische Studien*, 16/17 (1982): 283-312.

retrospect that the solution I had in mind can probably not be worked out in a satisfactory way.

(iv) *Consequentialism is groundless.* There is simply no evidence for it. You have evidence for a given theory if it is an essential component in the best explanation of some fairly uncontroversial facts. But consequentialism does not seem to contribute to the explanation of any facts at all. In a sense, consequentialism may of course be said to explain its own consequences—i.e. facts to the effect that certain particular actions are right and wrong, respectively— but such consequences can hardly be regarded as uncontroversial facts, and even if some of them are uncontroversial facts it would be a very controversial question whether consequentialism is part of the best explanation of them.

Moreover, a person's acceptance of consequentialism cannot be justified by the method of reflective equilibrium, since this presupposes that the person has justified nonmoral beliefs of a very complicated kind about the long-term consequences of alternative actions. It seems clear that we have no justified beliefs of this kind.<sup>19</sup> Hence, it is hard to see how a person's belief in consequentialism could be justified.

(v) *Consequentialism is useless in practice.* The point just mentioned—namely that we cannot have justified beliefs about the long-term consequences of alternative actions—also means that consequentialism cannot be applied in practice. We can never know that a given action is right or ought to be done according to the consequentialist principle. In some fairly extreme cases, we may have reason to believe that a given action is wrong, because it has foreseeable consequences which are so terrible that they are unlikely to be outweighed by unforeseeable good consequences or by the bad consequences of available alternatives. But this is not much help in ordinary life, where such terrible consequences are comparatively rare.

Of course, many or most proponents of objective consequentialism would claim that this is no objection at all. They would agree that consequentialism is not applicable in practice, but they would insist that this is irrelevant. They would say that all that can be asked of a theory of rightness is that it is *true*. A theory of rightness is not *meant* to be applicable in practice. What should be applicable in practice is something very different, namely a decision procedure or method of deliberation. In my view, this defence of consequentialism is unsatisfactory. Let us consider the point of practical inapplicability in more detail.

### **3. The inapplicability of objective consequentialism.**

Suppose I have nothing in particular to do this afternoon, and that I am thinking about going to the movies. There is a film I would very much like to see, and I have sufficient money in my pocket to buy a ticket. I am just about to make a decision to go, when I suddenly realize

---

<sup>19</sup>See Folke Tersman, *Reflective Equilibrium: An Essay in Moral Epistemology* (Stockholm: Almqvist&Wiksell, 1993), p.122; or Folke Tersman, "Utilitarianism and the Idea of Reflective Equilibrium", *Southern Journal of Philosophy*, 29 (1991): 395-406.

that I could use the money in a better way. I could use it to buy a ticket in a lottery where I can win a million dollars, which I can then donate to Oxfam. Surely, the consequences would be better if I donate a million dollars to Oxfam than if I go to the movie. Let us imagine that further alternatives can be neglected. Then according to consequentialism, I ought to buy the winning lottery ticket and send the money to Oxfam.

A natural objection to this way of reasoning is that I *cannot* buy the winning ticket, and hence, since "ought" implies "can", it is not the case that I ought to do so. However, it is not so clear that I cannot buy the winning ticket. Let us make the following further assumptions. The lottery is such that the tickets are numbered, and the buyer is allowed to specify the number of any particular ticket he or she wants to buy. There are ten thousand tickets which have not yet been sold. Among them is ticket #4711, which is in fact the winning ticket. There is no doubt that I can ask for ticket #4711. If I try to do so, I will in fact succeed. And if I ask for this ticket, I will win a million dollars, which I can then donate to Oxfam. It certainly seems that I *can* do this. Therefore, according to consequentialism, this is what I ought to do.

The problem, of course, is that I do not *know*, and *cannot* know, that ticket #4711 is the winning ticket. But why should this be a problem? As we have seen, objective consequentialism is a theory of rightness, not a decision procedure, and the important thing about a theory of rightness is that it gives true answers to moral questions; it is completely irrelevant whether we are able to know what these answers are. This is the very point of the distinction between theories of rightness and decision procedures!

A further feature of my example should be noticed. We can be pretty sure that it would be *wrong* for me to go to the movie, according to objective consequentialism. Surely, the consequences of this action are not as good as the consequences of my buying the winning ticket and sending a million dollars to Oxfam. (I may even keep a few dollars and see the film later.) And although I do not know *how* to buy the winning ticket, it is clear that this is something I can do. Therefore, going to the movie is not an optimal alternative. Therefore, it is morally wrong according to objective consequentialism. Notice, also, that this is a feature that my example seems to share with many situations in real life. In most cases, we may have fairly good reasons to believe that a given action is not optimal, but it is usually impossible to identify an optimal alternative in a given situation. Perhaps there is almost always the possibility of buying a winning ticket or backing a winning horse (or something like that) instead of doing what we actually do or contemplate.

In any case, the discussion above strongly suggests that objective consequentialism is unacceptable. It cannot be true that I ought to buy the winning ticket. A norm like that has no practical relevance whatsoever; it cannot be action guiding. I submit that moral norms should have practical relevance. That's what they are for. It also seems that this principle is what motivates the generally accepted slogan that "ought" implies "can".

Much the same point can be summarized as follows. Objective consequentialism says that you ought to maximize value. "Ought" implies "can", and "can" implies "knows how". Therefore, objective consequentialism implies that you know how to maximize value. But you don't know how to maximize value. Therefore, objective consequentialism is false. From a moral point of view, it is not sufficient that a given action, which you can do, has a morally desirable property. If you ought to perform the action which has this property, you must also be able to find out *which* action this is.

In view of this, it may be tempting to say that "ought" implies, not only "can", but also "can know that one ought". This would, however, be too strong. For one thing, we should not assume, without sufficient reason, that there is such a thing as moral knowledge. Rather, I suggest that "ought" (and "right") implies "there is some property, which this action has, and which makes this action morally obligatory (or morally right), and which is such that it is possible for the agent to know or to have good reasons to believe that the action has this property". In my opinion, this principle is just as plausible as the more well-known slogan that "ought" implies "can". But if this principle is true, objective consequentialism is false.<sup>20</sup> Q.E.D.

At this point, it might be objected that objective consequentialism may still be true for the *objective* sense of "ought", "right", and "wrong", even if it is false for the *subjective* sense of these terms. The distinction can be explained by saying that

an act is wrong in the objective sense if it is wrong in light of all the facts, knowable and unknowable, whereas it is wrong in the subjective sense if it is wrong in light of what the agent had good reason to believe.<sup>21</sup>

Allan Gibbard has claimed that the objective sense is not much use, "since what we need for moral guidance is criteria for what kinds of acts are wrong in the subjective sense".<sup>22</sup> The objective consequentialist may reply that this complaint is beside the point, since objective consequentialism is not meant to give moral guidance in the first place. However, it is far from clear that objective consequentialism can be true even for the objective sense of

---

<sup>20</sup>Besides, even if it could be known that a given action has better consequences than its alternatives, a sensible consequentialist should acknowledge that it may not be morally right. For its optimality may depend on the fact that the agent acts in a certain way later on, where a better alternative would have been another *combination* of actions now and later; see e.g. Lars Bergström, "Utilitarianism and future mistakes", *Theoria*, 43 (1977): 84-102.

<sup>21</sup>Allan Gibbard, *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Oxford: Clarendon Press, 1990), p. 42. For a very similar distinction, see Richard B. Brandt, *Ethical Theory* (Englewood Cliffs: Prentice-Hall, 1959), pp. 362-5).

<sup>22</sup>*Wise Choices, Apt Feelings*, p. 43. In my view, it is very doubtful that the subjective sense is what we need for moral guidance. When I want to find out what I ought to do in a given situation of choice, I want to know what I ought to do *in the situation as it actually is*, i.e. "in the light of all the facts". It seems that "what I ought to do in the light of what I have good reason to believe" may not be identical with what I ought to do in the situation as it actually is, and, if so, I am not primarily interested in what I ought to do in the subjective sense.

"ought". As far as I can see, it can hardly be true that we ought to maximize value "in light of all the facts, knowable and unknowable". For among all the facts is the fact that we can never know what will maximize value. And, in the light of this fact, it can hardly be the case that we ought (in the objective sense) to maximize value.

#### **4. Consequentialism in ordinary life.**

For the reasons given above, I believe that objective consequentialism has to be given up, even if it is still attractive in certain ways. But consequentialism—in a different sense—can still play an important role in practice. In moral decision making there is plenty of room for rules and strategies of a broadly consequentialist kind.

For example, I may apply the rule that the *foreseeable* consequences of a certain action are better than the *foreseeable* consequences of each alternative, then it ought to be done. This rule is clearly consequentialist, and in many situations it is a very reasonable rule. But it is not reasonable in every situation. Some foreseeable consequences may be more probable than others, and such differences may be important in some situations. Again, the foreseeable consequences of A may be better than those of B, but if there is also a small risk that A will lead to disaster, the agent ought perhaps to do B nevertheless.

In general, each one of the usual decision rules discussed in decision theory—such as minimax, minimax regret, principle of insufficient reason, dominance, maximization of expected utility, and so on<sup>23</sup>—has its use, but each one of them has to be qualified by a *ceteris paribus* clause; they hold, not without exception, but only "other things being equal". Each rule is unsuitable, from a moral perspective, in certain situations. Moreover, each rule presupposes a problem specification—usually in the form of a matrix of real numbers—where the relevant alternatives, possible states of nature, outcomes, values, and probabilities are described. A major problem for a decision maker is to determine the relevant matrix, but decision theory is not of much help here. So, from a moral point of view, each rule has to be supplemented with the proviso that the matrix to which the rule is applied is morally relevant. This too is covered by the *ceteris paribus* clause.

According to consequentialism the moral status of actions is determined by the value of outcomes or consequences. The value of consequences is in turn supposed to be determined by their more empirical features, as e.g. the extent to which they involve pleasant experiences or the satisfaction of needs, preferences, or interests. But it is difficult, if not impossible, to formulate a precise and generally acceptable rule according to which value is determined by empirical features. The rules which may be applied here need *ceteris paribus* clauses, which can also be made explicit in the general consequentialist rule itself. Another, essentially equivalent, possibility is to replace the general rule by various more specific rules, which mention particular empirical features. For example, instead of the general rule

---

<sup>23</sup>See e.g. Resnik, *Choices*.

that an action ought to be done if it maximizes value, we could have more specific rules saying that an action ought to be done, *ceteris paribus*, if it maximizes the empirical feature E— where E could be pleasure, the reduction of suffering, the satisfaction of preferences, the satisfaction of needs, and so on. Clearly, the *ceteris paribus* clause would be absolutely essential here.

Once we think of consequentialist rules as being qualified by *ceteris paribus* clauses, it is natural to allow further moral rules, which may not normally be thought of as consequentialist at all. An example may be the rule that people's rights ought not to be violated (*ceteris paribus*).<sup>24</sup> Other examples may be the usual rules that promises ought to be kept, parents ought to care for their children, people ought not to be punished for crimes they have not committed, and so on.

So which are the relevant moral rules? Well, as far as I can see, the relevant moral rules are those which are true, and I suggest that most or all of the rules indicated above are indeed true—as long as they are qualified by a *ceteris paribus* clause. Most of us would accept these rules. These are the rules we have learnt to accept. Of course, they can often point in different directions in actual situations of choice, but we have also had some instruction in how to handle such conflicts. We have learnt that some rules have more weight than others, for example.

## **5. Moral rules and theories of rightness.**

It should be emphasized that the moral rules I have in mind here are meant to be applicable in practice. Therefore, it may be thought that they belong to a decision procedure or method of deliberation rather than a theory of rightness. But this very distinction becomes open to doubt, when consequentialist rules are recognized as applicable in actual decision making. Is there any reason to assume that the moral rules cannot constitute a theory of rightness by themselves? Must there be some *other* theory of rightness, in addition to the moral rules we apply or should apply in practice? And if so, what is the relation between such a theory and the moral rules?

One possible reason for saying that the moral rules do not by themselves constitute a theory of rightness is that they do not determine rightness and wrongness in a unique way. The point is not that the rules may give rise to different classifications when they are used by different people, for this can easily happen with objective consequentialism as well. Rather, if objective consequentialism is a theory of rightness, but the conjunction of moral rules is not, it is probably because objective consequentialism identifies an certain right-making property which is (taken to be) common and specific to all morally right actions.

---

<sup>24</sup>Such rules may of course be presented in a clumsier but more consequentialist form. For example: a given action ought not to be done, *ceteris paribus*, if it leads to a state in which people's rights have been violated.

I agree that a simple conjunction of moral rules may not identify such a property. However, if we add to such a conjunction a general rule to the effect that an action is morally right if and only if it satisfies the (other) rules in the conjunction to a maximal degree, then we have explicitly introduced a property which is (taken to be) a right-making property which is also common and specific to all morally right actions. Moreover, I submit that some general rule of this kind is indeed part of ordinary morality.

It may now be objected that the property of maximizing the satisfaction of moral rules is not sufficiently precise or intersubjectively recognizable to play a fundamental role in a theory of rightness. But the same can be said of the property of maximizing value or promoting the good to a maximal degree. One may then try to replace this evaluative property by a right-making characteristic which is purely empirical, in the sense that it can "in principle" be detected by means of the ordinary empirical methods of science. For example, one might try to invent scientifically respectable and morally relevant explications of formulations such as "maximizes pleasure" or "maximizes the satisfaction of preferences". This task would be really difficult, partly because of the problem of interpersonal comparisons, but perhaps it is not impossible. Notice, however, that if it is possible, it may also be possible to invent a scientifically respectable explication of "satisfies the moral rules to a maximal degree". And then a conjunction of rules may determine rightness and wrongness in a unique way.

However, it is not clear to me that there *must exist* a purely empirical right-making property which is common and specific to all morally right actions. If this is required by a true theory of rightness, there may be no true theory of rightness. (Analogously, there may be no purely empirical property which is *truth*-making as well as common and specific to all *true* utterances.) Maybe, there is only a so-called family resemblance among morally right actions. If so, it seems that a theory of rightness would have to consist of a conjunction of fairly heterogenous moral rules.

Suppose that a conjunction of moral rules can constitute a moral theory. It must be admitted that such a theory is not very elegant. From a theoretical point of view, it would be nice if we could derive all true moral rules from some simple moral law—such as objective consequentialism—in connection with true empirical generalizations. However, I can see no reason at all to believe that this is possible. We have really no reason to believe that the applications of our rules usually or often pick out actions as right which are indeed right according to objective consequentialism. On the contrary, the usual counter-arguments against consequentialism suggest that there is very often a conflict between consequentialism and ordinary morality.

A more plausible idea is that our ordinary moral rules have a certain positive acceptance-value.<sup>25</sup> In other words, good consequences may flow from a general practice where people try to conform to these rules. That may be a reason for sticking to these rules, but this has nothing to do with objective consequentialism. Notice, also, that we have no reason to believe that our own particular system of moral rules has maximum acceptance-value. Our rules may be the result of an experimental selection process during the past history of mankind, but we cannot expect this process to have delivered the best possible rules already. Evolution has not had enough time for that, and human societies have often been too conservative to allow moral experimentation to a sufficient degree.

It is sometimes suggested that something like objective consequentialism is needed to solve the conflicts which arise between ordinary moral rules. For example, Richard Hare claims that

however well equipped we are with these relatively simple, prima facie, intuitive principles or dispositions [these are more or less what I have here called "moral rules"], we are bound to find ourselves in situations in which they conflict and in which, therefore, some other, non-intuitive kind of thinking is called for, to resolve the conflict.<sup>26</sup>

According to Hare, this non-intuitive, "critical" thinking involves "judgements which are the same as a careful act-utilitarian would make".<sup>27</sup> However, since objective consequentialism cannot be applied in practice, the only critical or quasi-critical thinking human beings are capable of would seem to be in terms of various consequentialist rules with *ceteris paribus* clauses (as exemplified above). For the purposes of conflict resolution, therefore, we are still restricted to "intuitive" moral rules.

But perhaps some version of objective consequentialism is needed to establish a *fact of the matter* as to whether given actions are morally right or wrong? Our ordinary moral rules may not be sufficient for that. At least, it seems that different people may come to different conclusions when they apply the same rules in a given case. However, if this is what a theory of rightness is for, and if, therefore, there is a fact of the matter as to which actions are right *independently* of our system of moral rules, then it seems that our rules may very well be completely mistaken. Therefore, the idea that objective consequentialism determines the moral facts is not easily combined with the acceptance of our ordinary moral rules. Furthermore, as I said in section 2 point (iv) above, we have really no reason to believe that objective consequentialism determines the moral facts.

---

<sup>25</sup>For the notion of acceptance-value or acceptance-utility, see e.g. David Lyons, *Forms and Limits of Utilitarianism* (Oxford: Clarendon Press, 1965), p. 140.

<sup>26</sup>See *Moral Thinking*, p. 40.

<sup>27</sup>Op.cit., p. 43.

## 6. Objective and subjective consequentialism.

However, someone might claim that objective consequentialism provides some kind of explanation of subjective consequentialism, which in turn explains many particular moral judgments which we actually make. In this way, we may seem to have a coherentist justification for objective consequentialism. For example, Torbjörn Tännsjö seems to reason along some such lines.<sup>28</sup> If he is right, objective consequentialism would have an important role to play after all. But I believe Tännsjö is mistaken.

I shall not say much about the second half of the justificatory argument—i.e. the claim that subjective consequentialism explains immediate normative judgements we make in particular cases—for this seems to be an empirical claim which would be very difficult to justify and which Tännsjö does not really make in any case. He merely indicates that this could be part of a coherentist justification of objective consequentialism.<sup>29</sup> Let me just point out that it would be really surprising if the claim were true, in view of the fact that consequentialist reasoning is often supposed to lead to results which are incompatible with our more intuitive judgements.

For my purposes, the more interesting part of Tännsjö's argument is the claim that objective consequentialism explains (or supports) subjective consequentialism. I am not absolutely sure that Tännsjö really makes this claim either. In any case, what he says is this:

If utilitarianism is correct, then, at least for a person who wants to be moral, it would be rational in many cases to try to maximize expected utility.<sup>30</sup>

Utilitarianism, in this context, is a hedonistic form of objective consequentialism, and presumably the recommendation to try to maximize expected utility (in many cases) is a (hedonistic) form of subjective consequentialism.<sup>31</sup> Clearly, however, what Tännsjö says here is highly controversial. In the first place, what is rational for a person depends primarily upon what he or she believes and desires. It can hardly depend upon whether utilitarianism is correct. But perhaps Tännsjö assumes that the person in question *believes* that utilitarianism

---

<sup>28</sup>See Torbjörn Tännsjö, "In Defence of Theory in Ethics", *Canadian Journal of Philosophy*, 25, 4 (December 1995): 571-94, pp. 586-91.

<sup>29</sup>What Tännsjö says is this: "If, in a particular case, we realize that a particular action does maximize expected utility, and if, in the same situation, we form the immediate judgment, without any conscious reasoning having taken place, that, from a moral point of view, this action is reasonable to perform, and if we conclude that it is reasonable to perform because it maximizes expected utility, this may constitute evidence, directly, for the maximizing method; and, indirectly, for the utilitarian formula. The utilitarian formula explains why it is reasonable to maximize expected utility (in these circumstances, for a person who wants to behave morally)", *ibid.*, p. 589.

<sup>30</sup>*Ibid.* Italics removed. See also the quotation in note 29.

<sup>31</sup>Tännsjö explains in a footnote that the relevant probabilities are (the agent's?) *subjective* probabilities. I assume, further, that the relevant utility function measures the agent's beliefs concerning the overall hedonistic value, for everyone affected by the action, in various possible states of affairs.

is correct. This would make the claim sound a bit more plausible, but the correctness or truth of utilitarianism would then be completely irrelevant. Besides, and more importantly, if utilitarianism is correct (true) it may often or even always be *morally wrong* to try to maximize expected utility. This depends, of course, upon whether some alternative would have better consequences—and this possibility can never be excluded. Now, presumably "a person who wants to be moral" wants to do what is morally right. How can it be "rational" for such a person to do what is morally wrong? Or is Tännsjö merely claiming that it would be rational for a person to try to maximize expected utility in those situations (if any) in which this would be the morally right thing to do? This is not much help.

Tännsjö seems to believe that "if, consistently, we try to maximize expected utility, then, in the long run, probably, we end up with better results than by consistently applying any conceivable alternative strategy".<sup>32</sup> But he admits that there exists no good argument for this belief.<sup>33</sup> And even if the belief were true, it does not follow that the maximizing method would pick out actions which are morally right (according to objective consequentialism) more often than not, or more often than would some alternative method. However, Tännsjö seems to have a different argument for adopting the maximizing method. He says that by using this method

we know that we concentrate on aspects that, according to utilitarianism, are of moral importance. And we know that we concentrate on no aspects other than these. This means that, by adopting the maximizing method, we are at least doing our best. *This* constitutes, for people who adhere to the utilitarian formula, the rationale behind the use of the maximizing method.<sup>34</sup>

This is not convincing. For example, subjective probability plays a central role in the maximizing method, but this is not an aspect that is of moral importance according to objective consequentialism ("the utilitarian formula"). And it is not easy to see in what sense "we are doing our best" by adopting the maximizing method. It is quite possible that there exists some alternative with better consequences. And it is quite possible that we apply the maximizing method to a decision matrix which leaves out some relevant information which we could have obtained if we had spent more time and money on empirical research concerning available alternatives and their consequences. So in what sense are we "doing our best"?

In general, then, an objective consequentialist has no reason to accept subjective consequentialism. This is something objective consequentialists have often pointed out.

---

<sup>32</sup>Op. cit., p. 589. He says that this is "probably true (of many utilitarians, at any rate)", if "probably" is taken in a subjective sense; see p. 590.

<sup>33</sup>See op. cit. p. 590.

<sup>34</sup>Ibid.

Conversely, subjective consequentialists have no reason to accept objective consequentialism. In particular, objective consequentialism is not an explanation—and even less the best explanation—of subjective consequentialism. On the other hand, some form of subjective consequentialism—possibly restricted by some *ceteris paribus* clause and supplemented by some further, less consequentialist rules—should be accepted (see section 4 above).

### **7. Morality without foundations.**

If the moral rules (or moral beliefs) we apply in practice are not supported by a more basic theory of rightness—such as objective consequentialism—it seems that they have no foundation at all. But why should we accept them, if they have no foundation?

We may distinguish between two different questions here, corresponding to the old distinction between decision methods and theories of rightness. First, why should we try to apply the moral rules, when we want to find out what we ought to do? The answer to this question is, I suggest, that the rules are (mainly) true. Second, why should we believe that the rules are (mainly) true, if they have no foundation? The answer to this question is that we *do* believe, of each one of them, that it is true. Moreover, as far as we can see at the moment, we have no reason to believe, of any one of them, that it is false. It is not irrational to stick to a moral belief when you have no reason to reject it. The fact, if it is a fact, that you lack a positive reason to accept it is not a reason to reject it.

Some beliefs are such that it would be irrational to stick to them without reason. Presumably, this is true of most scientific beliefs and of many specific everyday beliefs. Therefore, such beliefs can be expected to be based on reasons in most cases. For example, it would surely be irrational for me to believe that Richard Hare is at this very moment in London, if I have absolutely no positive reason to think so. But certain other beliefs can be rational even if they are not supported by reasons. Observational and introspective beliefs, for example, may belong here. Similarly for certain philosophical beliefs, such as the belief that some beliefs are rational as long as the believer has no reason to believe that they are false. I suggest that ordinary moral beliefs also belong to this latter category. If I am right, it may be quite rational to accept moral rules for which one cannot specify any foundation.

Of course, what we believe must be consistent. Hence, if the moral rules seem to have implications, in particular cases, which we find it hard to accept, we should try to modify our moral beliefs in such a way that they are consistent. Similarly, if we have reason to believe that our moral rules have certain characteristics which we believe that true moral rules do not have, then we have reason to modify our rules or our beliefs in such a way that what we believe is consistent. For example, I suppose most of us believe that true moral rules have a positive acceptance value (defined in terms of the satisfaction of basic human needs, say). If we find evidence which suggests that a certain moral rule has negative acceptance value, we

may therefore have reason to believe that the rule in question is not true after all. We may then reject the rule and thereby avoid inconsistency.

Consistency is not the same as coherence, and coherence is what gives justification. Therefore, it may also be rational to try to increase coherence by inventing more general moral rules which we can accept and by which the usual moral rules can be explained. In this way, we may be able go beyond mere consistency. However, it is possible that the degree of coherence that can be achieved for moral beliefs will always be much lower than the degree of coherence that exists for scientific beliefs.

In any case, it seems to me that a true theory of rightness may very well consist of a set of moral rules, with *ceteris paribus* clauses, which are not derived from a more basic moral theory and which therefore have no foundation. After all, we have no choice but to start with the moral rules we actually accept. We believe that these rules are true and that they should be applied in practice, and we have no reason to change this belief until we are faced with criticism which is sufficiently serious to require some kind of belief revision. Moreover, such revision has to be piecemeal in the style suggested by Neurath's metaphor of the sailors who must rebuild their ship on the open sea. Various consequentialist rules are part of our moral ship, but there are also important parts which are not obviously consequentialist, and the ship neither contains nor rests upon objective consequentialism. (Neither, of course, does it contain the principle that a moral rule is true if and only if, and because, we believe that it is true. This principle is not part of what we believe.)<sup>35</sup>

---

<sup>35</sup>An earlier version of this paper was presented at a conference at the University of Umeå in 1995. I am grateful for comments on this second version by Jonathan Dancy, Hans Mathlein, Christian Munthe, Wlodek Rabinowicz, Michael Slote, Folke Tersman, and Torbjörn Tännsjö.