

TWO FORMS OF CONSEQUENTIALISM, OR WHAT SHOULD WE EXPECT FROM A NORMATIVE THEORY?*

Andreas Lind
Lund University

Abstract

In this paper I argue that a version of consequentialism that is construed in terms of the maximization of *expected* value is preferable to a version construed in terms of the maximization of *actual* value. It is argued that only on the former account can we have a proper balance between the theoretical explanation of overall moral obligation and the idea that a normative theory should be practical.

Keywords

Objective consequentialism, decision-theoretic consequentialism, moral obligation, criterion of rightness, Frank Jackson.

1.

Contemporary act-consequentialism is divided on the question of how the *criterion* of rightness for an act or a course of action should be interpreted. According to *objective consequentialism* (OC), an act is right if and only if it would maximize actual value among those acts that are available to the agent. This version of the criterion has been defended by, among others, H. Sidgwick (1907), G. E. Moore (1903, 1912), P. Railton (1984), D. Parfit (1984) and D. Brink (1986). According to the decision-theoretic account of consequentialism (DTC), an act is right if and only if it maximizes expected value according to the agent's probability function at the time of action. This version has been defended by, among others, R. Brandt (1959), J. J. C. Smart (1973), F. Jackson (1991, 1997) and M. Zimmerman (forthcoming). *If we believe that act-consequential-*

* Earlier versions of this paper were presented at the Higher Seminar in practical philosophy at Lund University, Sweden, at the Analytical Philosophy and European Culture conference in Genova, Italy and at the 2nd National Meeting for Analytic Philosophy in Porto, Portugal. For many valuable comments, I would like to thank the participants at these presentations. I am especially indebted to comments by Erik Olsson, David Alm, Daniel Svensson, Johan Brännmark, Caj Strandberg and Björn Petersson. Special thanks to Wlodek Rabinowicz.

ism as a generic normative theory is correct, *how* do we decide whether we ought to embrace OC or DTC?¹ I happen to believe that DTC is the best version of consequentialism. In my opinion it is the only version of the two that has the prospects of succeeding as a genuinely normative theory in ethics. It can do so, I will argue, because its account of *overall moral obligation* is properly balanced by, on the one hand, a plausible theoretical explanation of the notions of ‘ought’ and ‘right’ and, on the other, a plausible picture of a normative theory’s practical aspects.

OC and DTC cannot both be right since, as we will see, they (sometimes) give conflicting prescriptions as to what an agent ought to do. Prima facie, however, we have reasons to believe that ‘ought’ have senses that pull in opposite directions – one supporting OC and one supporting DTC.² If that is true, one might think that the two criteria need not be in conflict with each other.³ Indeed, ‘ought’ is obviously ambiguous when we take into account *all* of its different senses but it is *not* so, I will argue, regarding the relevant sense under discussion here, namely, the sense that express overall moral obligation.

I will also argue that using OC as a criterion of rightness and DTC as its complementary decision procedure is untenable (section 3.). If OC and DTC cannot coexist or be used in a complementary way, the only option left is to argue that one version is preferable to the other (which in effect means that only one version captures the relevant sense of ‘ought’). That is my aim in this paper.

I will draw heavily on Frank Jackson’s version of DTC. I consider it to be the soundest and most fruitful version of consequentialism that we have at present. One of Jackson’s strategies is to argue against OC via examples. He shows, forcefully to my mind, that OC gives intuitively wrong answers in certain real-life cases. The upshot of that strategy is that if we take OC’s version of ought, oughtobjective as I will later call it, as the relevant sense of overall moral obligation, we will end up with highly undesirable results. In section 2 I state Jackson’s view of how we should regard DTC together with my own view of this matter. In section 4 I discuss and acknowledge an argument launched by Fred Feldman to the effect that DTC might not be so practical after all.

There are some things I will not discuss here. For instance, the question of how we should understand *what* it is that should be maximized according to OC and DTC, that is, questions of value. I will also simply assume the maximizing approach, though I think my findings can be generalized to satisficing approaches as well. Finally, I will not discuss objections traditionally brought against consequentialism, such as, for instance, whether it is an excessively demanding theory.

2.

I suggested above that one of the reasons why it is hard to take a stand on the question whether we should embrace OC or DTC is that they appear to focus on different senses of ‘ought’ and ‘right’. Another reason is that, regardless of where our sympathies might lie at the outset, the aspects thought to be the greatest strengths of these theories are also their greatest weaknesses. I will start with OC and state its rightness criterion in a little more precise way:

(1) It is right for S to do A iff no total state of affairs that would be a consequence of S’s doing any alternative to A would be better than the (actual) total state of affairs that would be a consequence of S’s doing A.⁴

¹ The conditional is meant to safeguard the present discussion from yet more general questions such as whether we have any reasons to believe that act-consequentialism is the correct normative theory in the first place. An implication of this is that I am not going to compare OC and DTC with other normative theories.

² Cf. A. C. Ewing (1947), chapter 4

³ Cf. Jackson (1991: 472) and David Sosa (1993: 108f.).

⁴ Cf. Sosa (1993: 101). Is an action wrong if it is not right? Bentham and Mill seem to have acknowledged a sort of

Let us call a right act thus defined as $\text{right}_{\text{objective}}$. If an agent ought to do the act that is $\text{right}_{\text{objective}}$, let us call this $\text{ought}_{\text{objective}}$. OC countenances the idea that whether an act is right or wrong is a question about that act's objective status. The (possible) gap between our beliefs about an act, say, that it is the right act to do in a particular situation and the fact that the act really is right can be explained in a quite straightforward way by OC: the act is right if it maximizes value and that is so regardless of what we believe about it. Adding qualifiers to the agent's epistemic state does not, on the face of it, make X the right thing to do. Even if we say something like "She was justified in believing that X was the right thing to do" or "It was reasonable for her to believe...", we still have to acknowledge, the argument goes, a gap here. Sometimes this proposed version of objectivity is supported by a metaphysical assumption to the effect that for every possible action it is always *a fact of the matter* what outcome that action would have if it were performed.⁵ I will, however, not let my evaluation of OC depend on whether we accept or reject this debatable assumption.

But the standard objection against OC is a sceptical one: the information required to know whether an act maximizes value is nowhere to be had. So even if the metaphysical assumption is true, we can simply never know or even be fairly certain that an act is right and thus ought to be done according to OC.⁶

We can in contrast formulate DTC's rightness criterion roughly like this:

(2) It is right for S to do A iff there is no other act, of all the available alternatives to A, with a higher expected value according to S's probability function at the time of action.⁷

DTC is explicit in its supposition that a normative theory should be functional in practice and is thus taking seriously the idea that ought implies can, not just in the sense that the agent can do what she ought to do but also in the sense that she can know what she ought to do.⁸ In establishing the rightness of an act via a subjective probability function, DTC is bringing the *agent* directly into the picture. In this way the information needed to ascertain whether an act is right and ought to be done is actually and not only ideally within the reach of the acting agent. Let us call a right act as defined by DTC, $\text{right}_{\text{subjective}}$.⁹ If an agent ought to do the act that is $\text{right}_{\text{subjective}}$, let us

continuum in which actions can be more or less 'right'. Mill said "that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness". John Stuart Mill (1998: 55, section 2:2). With Sidgwick, and more clearly so with Moore, this idea is abandoned. If an action is not right, according to Moore, then it is wrong. I will use the latter conception below.

⁵ Cf. Erik Carlson (1995: 20). Carlson calls this assumption 'counterfactual determinism' and he concedes that it is debatable. If that assumption is false, he agrees, some form of (objective) probabilism would be more plausible. Let me also stress that there is a further distinction to be made between 'right' on the one hand and 'ought' to do and 'duty' on the other, as emphasized by Moore in *Ethics*, chapter 1. If two or more actions produce equal amounts of value and that these actions produce more value than any other action open to the agent, then all these actions will be equally right. Actions that we 'ought' to do or that it is our 'duty' to do are actions that produce more value than any other action open to the agent. Hence, 'right' and 'ought' are not coextensive since an action can be right without being an action that the agent ought to. Every action that ought to be done is, however, also right. But the corresponding relationship does not hold between 'wrong' and ought not to be done, since every action that is wrong is an action that the agent ought not to do. Moore's discussion in *Ethics* is more extensive than I have conveyed here. My short presentation is just meant to serve as a reminder.

⁶ This was something Moore was very clear about. See chapter 5 in *Principia Ethica*

⁷ According to Jackson, "the rule of action is to maximize $iPr(Oi/Aj) \times V(Oi)$, where Pr is the agent's probability function at the time, V is consequentialism's value function, O_i are the possible outcomes, and A_j are the possible actions" (1991: 463f.) A short reminder: I am silent here, as is Jackson, on the question of what kind of value V represents. From now on all references to Jackson will be inserted continuously in the text.

⁸ Cf. Lars Bergström (1996: 83).

⁹ The terminology can be a bit confusing. As Railton defines it, "*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" and furthermore that "one uses and follows a distinctively

call this ought_{subjective}.¹⁰ Let me immediately stress that it is subjective probabilities that are at play here. This, however, does not preclude agents from making mistakes and thus that they can be wrong about the action that they ought to do. The possibility of wrongdoing on this account is therefore not precluded. So even if I speak of rights_{subjective} and oughts_{subjective} there is still a distinct possibility of error. I would therefore claim that DTC also is a version that accounts for the objectivity of actions being right or wrong. We must assume this since it would, of course, be wildly counterintuitive to let rightness and wrongness depend purely on what an agent believes about the situation without ensuring that these beliefs are reasonable or justified. DTC obviously does not imply the latter idea. I will briefly come back to this in section 3.

But the standard objection against DTC is still that the gap between the rightness of an act and what the agent believes about the act is not wide enough. DTC, the argument goes, is simply making rightness of acts into an affair of individual agents' beliefs, with an untenable relativity or subjectivity in its wake.

Before we proceed to address this objection, let us see how Jackson argues for DTC with a by now famous example about Jill the physician. One of the merits of DTC is that it allegedly gives the intuitively right answer about what Jill should do:

Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug. What should Jill do?

The possible outcomes we need to consider are: a complete cure for John, a partial cure, and death. It is clear how to rank them: a complete cure is best, followed by a partial cure, and worst is John's death. That is how Jill does, and also how she *ought* to, rank them. But how do we move from that ranking to a resolution concerning what Jill ought to do? The obvious answer is to take a leaf out of decision theory's book and take the results of multiplying the value of each possible outcome of each contemplated action by Jill's subjective probability of that outcome given that the action is performed, summing these for each action, and then designating the action with the greatest sum as what ought to be done. In our example there will be three sums to consider, namely:

$$\begin{array}{ll}
 Pr(\text{partial cure}/\text{drug A taken}) & \times V(\text{partial cure}) \\
 & + Pr(\text{no change}/\text{drug A taken}) \\
 & \times V(\text{no change}); \\
 Pr(\text{complete cure}/\text{drug B taken}) & \times V(\text{complete cure}) \\
 & + Pr(\text{death}/\text{drug B taken}) \\
 & \times V(\text{death}); \text{ and} \\
 Pr(\text{complete cure}/\text{drug C taken}) & \times V(\text{complete cure}) \\
 & + Pr(\text{death}/\text{drug C taken}) \\
 & \times V(\text{death})
 \end{array}$$

consequentialist mode of decision making, consciously aiming at the overall good and conscientiously using the best available information with the greatest possible vigour." (1984: 152). In *that* sense, DTC is not a version of subjective consequentialism since it is ordinarily not taken as a decision procedure nor does it say that we should try to maximize value (it is *expected* value that is to be maximized).

¹⁰ From now on I will use 'should' as neutral between 'ought_{objective}' and 'ought_{subjective}'. When I say, for instance, that 'and agent should do so and so' I hope it will be clear from the context that we have, say, intuitions that pull in a certain direction and that we have them whether or not we have taken a stand for OC or DTC.

Obviously, in the situation as described, the first will take the highest value, and so we get the answer that Jill should prescribe drug A. The obvious answer all along. (p. 462f.)

So DTC gives us the obvious and correct answer about what Jill should do – she should prescribe drug A, the drug that is likely to relieve the skin condition but not completely cure it. So the ‘should’ here is to be understood as ‘oughtsubjective’. ‘But a better ‘total state of affairs’ is presumably one where John is completely cured and so it seems that either OC would say that Jill should prescribe drug B or OC would say that Jill should prescribe drug C. We certainly balk at that implication. The lesson is that we should not do the act that maximizes actual value.¹¹

In a second version of the example Jill only has two drugs to prescribe. Drug X has a 90% chance of curing John but also a 10% chance of killing him; drug Y has a 50% chance of curing John but no bad side effects. “It is clear”, says Jackson, that Jill “should prescribe Y, and yet that course of action is not the course of action most likely to have the best results for it is not the course of action most likely to be objectively right” (p. 467.). The lesson is that we should not try to do the act that is most likely to maximize actual value.¹²

Hence the first part of the attack on OC is to claim that it gives the wrong answers in the drugs case. The next step is, according to Jackson, that “we have to see consequentialism as containing as a constitutive part prescriptions for action” (p. 466.) and “the fact that a course of action would have the best results is not in itself a guide to action, for a guide to action must in some appropriate sense be present to the agent’s mind” (p. 467.). Furthermore, “We need, if you like, a story from the inside of an agent to be part of any theory which is properly a theory in ethics, and having the best consequences is a story from the outside” (p. 467.). Let us pause a bit here and comment on something that might seem a bit odd in Jackson’s line of reasoning here. If he says that the best result, the act that is rightobjective, is not a guide to action, why does he bother to argue that OC gets the wrong results in the examples above? If something is not action-guiding, how can it lead to the wrong result? If anything, he should say that OC gets no result out of these examples. But that is perhaps too strong. Let us start by saying something very general (and simplified) about what it means for a normative theory to be action-guiding. Here is a question: Would we say that a principle is action-guiding if it is usable (I will explain this below) in practice for an agent, S, but every time S is guided by this principle it leads her, through no fault of her own, to the wrong result? There is something strange in saying that it is. If I write an instruction how to travel from place X to Y and everybody (including myself) who follows the instruction

¹¹ We are now in a position to state why the following version (sometimes described as a version of subjectivism) should be rejected:

(3) It is right for S to do A iff there is no other act, of all the available alternatives to A, such as it is most reasonable for S to expect that A would maximize actual value.

Jill’s evidence points clearly to the fact that either the prescription of drug B or the prescription of drug C would lead to maximization of actual value. In short, that (3) implies that it will lead to the same (unfortunate) prescription as (1). “[I]t is prescribing drug A which is the intuitively correct course of action for Jill despite the fact that she knows that it will not have the best consequences.” (p. 466. *Italics in original.*) As is clear from what I have said above, Jackson is a subjectivist in the sense of (2). Cf. “It is the subjective notion [of ‘ought’] – spelt out in terms of expected value – which is fundamental...”. (1997: 53).

¹² Jackson claims that with OC “the moral decision problem should be approached by setting oneself the goal of doing what is objectively right – the action that has in fact the best consequences – and then performing the action which the empirical evidence suggests is most likely to have this property” (p. 467.). What Jackson in effect is implying here is that any decision procedure compatible with OC’s criterion of rightness must have as a constitutive part that the agent sets herself a particular goal, namely, the goal of doing the objectively right thing. I think Jackson overreaches himself here. There are good (consequentialist) reasons for why an agent should not have that as a goal each time she acts. However, when an agent is trying to figure out what to do in a conscious and deliberate way (as in the example), and her evidence points in a direction that there actually is an act that maximizes value, then, it seems, the moral decision problem should be approached in the way Jackson says it should.

ends up in Z, then we could say either of two things about my instruction: (i) that it does not provide the agent with a definite instruction about what to do, given her information, or (ii) that it guides the agent to actions that are bad or inefficient in some sense. Consider the following condition, stated by Holly M. Smith:

(A) Agent S uses principle P as an *internal guide* for deciding to do act A if and only if S chooses to do A out of a desire to conform to P and a belief that A does conform (1988: 92).¹³

We can easily imagine an agent in the example above satisfying this condition. She wants to go to Y and she also has a desire to conform to my instruction and a belief that it does so conform. So my instruction is in this (internal) sense action-guiding. But my instruction *fails* to be action-guiding in another (external) sense:

(B) Agent S uses principle P as an *external guide* for deciding to do act A if and only if A conforms to P, and S does A out of a desire to conform to P and a belief that A does conform (Ibid.).

In the external sense, something went wrong when the agent ended up in Z and not, as intended, in Y. The reference to what S *does* in (B) implies that some results are supposed to issue from this external guide and if the results are not as stated, my instruction fails to be action-guiding. An agent can therefore satisfy (A) but fail to satisfy (B). I think this is also the reason why we above could say that my instruction is action-guiding, but that it is bad or inefficient.

A principle is ‘usable’, according to Smith, “just in case the agent is able...to use it in the sense of (A) and (B)” (Ibid.: 92). For the remainder of this discussion I will, however, ignore (B) and focus on (A).¹⁴ Jackson’s two ideas that a principle or a course of action ‘must in some appropriate sense be [able to be] present to the agent’s mind’ and that we need ‘a story from the inside of the agent’ suits the idea expressed in (A). This means that there must be some kind of appropriate fit between the principle and the agent. For the sake of simplicity I will here labour with an agent who is cognitively, psychologically and intellectually mature and normal. No further analysis of what this means will be carried out here. Suffice to say is that she is pretty much like you and me. The question we shall ask is whether an agent like this is able to use OC or DTC in a way that guides or could guide her actions.

When R. E. Bales made the sharp distinction between utilitarianism (consequentialism) taken as a criterion of rightness and taken as a decision procedure, he argued *against* the conception that accepting an “account of right-making characteristics somehow commits one *a priori* to a particular decision-making procedure” (1971: 263). In line with this Railton has argued that “it becomes an empirical question (though not an easy one) which modes of decision making should be employed and when” (1984: 156). A defender of DTC would not object to these characterizations. But hand in hand with the distinction between the rightness criterion and the decision procedure there is also a conception that the criterion can be a guide to action in on its own. Thus Sosa says that “The fact that a course of action would have the best results is indeed a guide to action, simply not one we are always in a position to use” (1993: 111). And Bales claimed that the act-utilitarian “has told us which alternative, under one description, to perform: the one which would maximize utility” (1971: 261). To say so is to say that something can be a guide to action despite the fact that the information needed to establish whether something in fact maximizes value cannot be had. *Pace* Sosa, I would say that this guide to action is one *we* are never in a posi-

¹³ Although I will not pursue the point here, it seems reasonable to demand that the belief that act A conforms to P must, in some sense, be justified. If an agent falsely believes that an act conforms to a principle, then it seems odd to say that she is guided by that principle in deciding to do a particular act.

¹⁴ I do this because it simply would be too strong to claim that for a principle to be action-guiding it must be so in both the internal and the external sense. Cf. Smith (1988: 94).

tion to use. But the question we must ask is how Bales, Sosa and others can say that OC has told *us* anything at all about which alternative to perform or that there is anything at all that deserves the name of being a guide to action on their account. If I am a normal agent in the way indicated above, how can I (justifiably) *believe* that an act conforms to the criterion (in the way indicated by (A)) when I *know* beforehand that I do not have, and cannot acquire, the considered necessary information? They seem to claim that because there is a fact of the matter, there will also be a sense in which this will be a fact of the matter *for me*. But that entailment seems only to follow on the assumption that there is a way for me to find out that it is a fact of the matter for me. It is like giving someone a manual to build a machine in a language which we know she does not master and then say “I, of course, know that you do not master the language in which this manual is written but if you want to build the machine it is only to follow the instructions”. The manual, we can imagine, is action-guiding for someone who masters its language, but it would be odd to say that it is action-guiding period. The difference between this example and OC’s account of rightobjective is that an agent can learn to master the language. But as in the movie *Contact* we can imagine that we get a message from some alien life form in our fax machine with a heading in English saying “Manual to build spaceship. Follow it if you want to build one.” with the rest of the text in an alien language that none of us know. In line with Smith’s condition (A) above, we can say that I might surely have a desire to build the spaceship with some help from the instruction but I surely cannot have a belief that I will be thus helped. Try to imagine Sosa’s and Bales’ stories with regard to this “action-guiding manual”. I would not know what to make of their claims. (I will argue against a possible answer in the next section.)

It might however be that there are many other areas in which there are ‘facts of the matter’ that are relevant for me, even if it is in principle impossible for me to find out about them. If, unbeknownst to me and everybody else, my favorite brand of whiskey in fact is a cause of cancer it is certainly hard to deny that this fact in some sense is relevant for me. If I knew that the whiskey was a cause of cancer I should of course stop drinking it. As things are now, however, when no evidence whatsoever points in the direction that it actually is a cause of cancer, we can still say that it would be a good thing for me if I did not drink the whiskey. But if something is a good thing for me, we might wonder, how can it be that I should not do it? Jackson’s answer is to admit that we have “an annoying profusion of ‘oughts’”:

I think that we have no alternative but to recognize a whole range of oughts – what she [the agent] ought to do by the lights of her beliefs at the time of action, what she ought to do by the lights of what she later establishes (a retrospective ought, as it is sometimes put), what she ought to do by the lights of one or another onlooker who has different information on the subject, and, what is more, what she ought to do by God’s lights, that is, by the lights of one who *knows* what will and would happen for each and every course of action. (p. 471 f.)

Consider the following oughts corresponding to this passage: ought_{subjective,t1}, ought_{subjective,t2}, ought_{2nd person}, ought_{God}. As I said above, the debate is currently structured around the fact that OC is trying to make ought_{objective} the pertinent sense of overall moral obligation and DTC is trying to make ought_{subjective} pertinent.¹⁵ Jackson says that for DTC “the ought most immediately relevant to action [is] the ought which [it is] the primary business of an ethical theory to deliver”. Ought_{subjective} is the relevant ought since an agent must use the information that is available to her at the time of action. In a footnote to the quoted section above, Jackson says “There are also the various nonmoral oughts – prudential etc., but that is another, and here irrelevant,

¹⁵ Gilbert Harman provides the beginnings of an inventory of different types of moral and non-moral oughts (1977: chapter 7). See also Ewing (1947).

dimension of variation.” This shows that Jackson takes all of the various oughts in the quoted section as moral oughts. But I believe that we must in fact re-center the debate a bit. If we agree that a normative theory, and *a fortiori* consequentialism, should contain as a constitutive part prescriptions for action, then it seems that the profusion of oughts is not so annoying after all. For we can agree that all the senses of ‘ought’ that Jackson mentions are legitimate, even that all of them are moral oughts, without agreeing that all of them are relevant to the normative domain associated with overall moral obligation. Prescriptions for actions are of no use (and perhaps the notion is even incoherent) if they are not capable to guide us, the acting agents. I would like to interpret Jackson’s claim that we need a story from the inside as the idea that action-guidingness is an agent-relative notion, that is, that it is relative to an agent’s beliefs and the evidence that is available to her at the time of action. Only an agent-relative notion of action-guidingness can answer the call for a normative theory to be practical. It is, to repeat, not to deny that there is a story from the outside, rather, it is to deny that that story is relevant to normative ethics.

To recap this last argument. There are many senses of ‘ought’; there are even many senses of what we have called the moral ‘ought’. It may or may not be that all of them have a certain valuable use within a specified domain. I have argued that for a normative theory only the practical oughtsubjective is appropriate. The purpose of that sense of ought is clear, namely, to support and reinsure the practical part of the normative theory of which it is a part. OC, on the other hand, focuses on oughtobjective, that is, the ought that is most appropriate for someone who has all the relevant information – a God-like being. My argument is, however, not sufficient as an argument that DTC is the correct version of consequentialism, since other versions of consequentialism may tell a better story of how we should regard oughtsubjective, but it points in a direction that if we must decide between OC and DTC, then DTC is the better option.

How do we acknowledge the fact that we believe that it would be a good thing for me not to drink my favorite whiskey? If I had God-like information I would presumably not drink it – In fact I oughtobjective not drink it. We *have* a conception of the oughtobjective, there is no need to deny that (although I believe it is not as clear as it is sometimes considered) and in having that conception we are led to the belief that I should not drink the whiskey. If the argument in this section is sound, however, that only goes to show that the purpose of the oughtobjective is not the purpose fit for a normative theory in ethics (unless, that is, we are discussing a normative theory for Gods). But what exactly is the use of oughtobjective then? I will take the easy way out here and say that it is not for the defender of DTC to specify a use or a purpose for it. One idea, however, is that it has a use in combination with DTC as a decision procedure. I will argue against that conception now.

3.

In this section I am going to argue for two things. First, that we should not accept a ‘package deal’,¹⁶ which is the suggestion that OC sets the criterion of rightness and DTC is its complementary decision procedure. The second argument is a response to a (possible) worry that DTC connects the deontic and the evaluative dimensions (explained below) too tightly. Corresponding to that response I will launch an argument to the effect that OC drives them too far apart. Let us begin with the package deal. Our working assumption has been that OC and DTC cannot both be right as criteria of rightness. A defender of the package deal will agree with this but claim that they are compatible as criterion and decision procedure. In the first place, then, an agent oughtobjective

¹⁶ I owe this label to Fred Feldman and the idea is discussed in his ‘Actual Utility, the Objection from Impracticability, and the Move to Expected Utility’, (forthcoming).

to do the act that is right_{objective}. But if she cannot figure out which act that is, she should follow a decision procedure which is, by hypothesis, defined by DTC and do the act that is right_{subjective}.¹⁷ The upshot is, again, that it is right_{objective} as defined by OC that really is the pertinent sense of overall moral obligation. An agent is successful in following the decision procedure in so far as she is doing the act that she ought_{subjective} to do. This idea can be coupled with a further argument for the package deal, namely, that it is not a matter of maximizing the number of right actions but to maximize whatever it is that has (intrinsic) value. An objection to this is of course the example with Jill. All her evidence positively indicates that prescribing drug A will not maximize actual value but still, on our preferred reading, what she ought_{subjective} to do.

The real problem with the package deal, as I see it, is that it reveals yet another and more serious problem for the acting agent who wonders about what to do. We have another profusion of oughts, a profusion that applies to the agent, from her own perspective, at the time of action. Let us assume that act A is right_{objective}. If she cannot identify A, she ought_{subjective} to do the act that is right_{subjective}. Let us assume that this is act B and that she has reached that conclusion by means of DTC. Let us also assume that she does not know that A is right_{objective} but that she knows that B is right_{subjective}. In that case she should do B – but B is *not* right_{objective} so her favored theory tells her to do the act that is not right. That seems rather strange. What we have, in effect, is an argument against OC as being action-guiding in referred to above. Prima facie it is only an argument against using a principle as an *external guide*, that is, (B). Still, there is something very awkward if we, so to speak, in principle are incapable in connecting (A) with (B). I will however not pursue this point any further here.

There is, however, a reply to the line of reasoning proposed above which turns on the distinction between what I will call the deontic dimension and the evaluative dimension.¹⁸ What the defender of the package deal can say is that the even if the agent in the previous example is not doing the act that is right_{objective} she is nevertheless not to be blamed for it. She followed the best decision procedure available to her at the time of action and in so doing she did something that led (say) to a great deal of value.¹⁹ But I believe that, except in highly unusual cases where an agent knows that a particular act is right_{objective}, the package deal is driving the evaluative dimension too far away from the deontic dimension. If doing the right_{objective} thing is what maximizes value and we are hardly ever in a position to know what maximizes value we will rarely, if ever, do the right_{objective} thing. We will therefore only be praiseworthy or blameworthy for succeeding or failing to do what is right_{subjective}, that is, what maximizes expected value. Assuming the package deal, there is no deontic ‘right’ or ‘wrong’ attached to that part. When I said above that the agent “ought_{subjective} to do B” that cannot as a result be an ‘ought’ that hooks up with the deontic ‘right’ in the right way. Let us instead call that the ought_{evaluative}, in order to make a difference between the evaluation of acts and the evaluation of agents. An agent that does what she ought_{evaluative} to do is on this suggestion without blame.

¹⁷ I am simplifying here. Maybe there is another and better decision procedure available to the agent. In that case she should obviously follow that instead of DTC. “Better” can here be understood in (at least) two ways: It may be better in the sense that in following it, the agent comes closer to doing the right_{objective} act or that it is better in that it has a preferable notion of right_{subjective}.

¹⁸ A word on the terminology here. I agree with Jonathan Dancy when he says that “Within the normative domain, considered as a whole, there are two potentially distinct parts. There is the evaluative, and there is the deontic. The evaluative is the realm of value, of good and bad, of evil and benign. The deontic is the realm of ‘oughts’ – of requirements, of demands, of right and wrong, and of reasons.” (2000: 29). I will nevertheless make use of a less nuanced terminology here, where I will treat the deontic as involving ‘right’ and ‘wrong’ and the evaluative as involving questions of ‘praise’ and ‘blame’. Nothing turns on this terminology here, however.

¹⁹ Again I follow Feldman in referring to Timmons’ *Moral Theory* for a line of defense along these lines.

With the package deal we presumably get the same answer as DTC in the drugs examples. According to DTC, Jill ought_{subjective} to prescribe drug *A* in the first and drug *Y* in the second example. In so doing she does the right_{subjective} thing (which is a deontic sense). According to the package deal she should do the same things – but only in the ought_{evaluative} sense. The problem with this proposal, and the thing that ultimately should lead us to reject it, is that the package deal is left only with the ought_{evaluative} without being able to connect it with the deontic sense. As I see it, it combines two bad features into an unworkable whole: It adopts OC's idea of a right_{objective} act (with all the problems stated above) and downgrades DTC into a decision procedure, and in so doing the practical side of the theory is left only with the ought_{evaluative}. But how do we praise or blame an agent if we are barred from connecting praise and blame from what is right or wrong? Well, we *could* say that the agent has made everything that we reasonably can expect of her. If so, however, it is far from clear what the function of OC's criterion of rightness is.

But if it is true that the package deal drives the deontic and the evaluative too far away there is a corresponding worry that DTC connects them too tightly. The argument I will use to respond to this worry can also be used to respond to the worry I mentioned at the beginning, namely, that DTC implies an untenable relativity since it makes the rightness of an act into an affair of individual agents' beliefs. I think Jackson has a perfectly good defense against *culpable ignorance* so that is not the main worry here. That defense turns on an old story often told by consequentialists, namely that we often have good consequentialist reasons to seek new information. When an agent faces a choice of, say, doing either *A* or *B* there will, at least in most cases, be a further act, *C*, which will be that of seeking more information about whether *A* or *B* will be right_{subjective}, and based on the new information decide whether *A* or *B* is right_{subjective}. The agent should first and foremost do *C* if that is right_{subjective}.²⁰ If *C* is right_{subjective} and she does not consider it, she is doing something that is wrong_{subjective} and should, presumably, be blamed for it. The worry is rather if the agent's beliefs are faulty from the beginning and so leads the agent to a distorted picture of the decision situation.²¹ Based on her initial beliefs she may not consider *C* as being right_{subjective} or perhaps even, if she does so, that she will acquire even more false beliefs if she seeks new information.

So an agent may consider act *A* to be right_{subjective} but because her beliefs are false from the start, *A* is actually wrong_{subjective}. If she had good reasons to believe that *A* is right_{subjective}, say, that she had done anything within her powers to check the evidence, made no mistakes in her calculations etc., then she is not blameworthy. In that sense there is a well-balanced distinction between the deontic and the evaluative dimension. The worry here is the implication that *A* actually turns out to be right_{subjective}, since it is a function of her beliefs and probability function. It is no good idea to suggest that the agent ought_{subjective} to do *A*, presumably because she is not blameworthy if she does it, but still hold that it is wrong_{subjective}. That will lead to the same conclusion as with OC (and the package deal) that we do not have the right connection between the evaluative and the deontic dimensions, because the ought_{subjective} will in that case only turn into an ought_{evaluative}.²²

²⁰ It may of course be that the act of seeking new evidence will not be the act that is right_{subjective}. In the example quoted above with Jill, the formulation that "Careful consideration of the literature has led Jill to the following opinions", is presumably a case where Jill's considerations have not led her to seek further information.

²¹ This, I take it, is the force behind Sosa's claim that "If it were *constitutive* of right action that its agent believed that act to be right, ignorance could not be the cause of a failure to act rightly." (1993: 112).

²² One could of course argue that as moral agents we should act so as to reach a certain level beyond which we are not morally reproachable. Whether we do what is right (objectively or subjectively) is not the real deal. What matters is that we are not doing something for which we can be blamed. Rightness is only a regulative ideal, perhaps. I think that there is something plausible in this approach but I am not sure whether it is very strong. I believe that what we want to do as conscientious agents is to do what is right and not merely what takes us beyond the of reach blameworthiness.

What we want, and indeed what we can have with DTC, is an account of why acting on false beliefs does not constitute right action. Consider Jackson's example with Jill once again. The formulation that "careful consideration of the literature" has led Jill to her opinions may be a possible solution to our problem. No one, I believe, would be convinced if Jackson had written something like "A quick and swift glance at the literature has led Jill...". I think that the original formulation was meant to imply that Jill had availed herself of the available evidence. I also think that we should follow Zimmerman²³ in saying that available evidence is evidence that an agent *ought* to avail herself. This is an *epistemic* 'ought'. It is not a moral ought since there are many cases in which we do not claim that an agent is morally blameworthy for not having considered available evidence. I cannot adequately explore the question of how we are to understand this epistemic ought here. I can only hint at some ways in which we do not want to understand it. What we do not want – on either of the accounts – is a strong or highly idealized notion of the epistemic ought. It will, for instance, not do to claim that an agent must avail herself of *all* the available information since that would obviously make all accounts impractical. Maybe we must invoke something like a threshold limit in order to ensure that acting agents can be practically and not only ideally within the reach of doing what they ought to do, that is doing what is right (whether on DTC or with the package deal). This does not imply that it would be an easy thing to what one ought to do. It is just that it must be *possible*, through a conscious and deliberate examination of the evidence, to act as we should. Much, much more needs to be said of the epistemic ought. Something I do not have the space for here.

4.

One serious objection against DTC is that it impractical! Feldman claims that the move to expected value is infeasible on this ground. Except in rare cases, he argues, the identification of the act that is rightsubjective is impossible. Remember that on Jackson's account the agent should assign a value to each possible outcome on each possible action and multiply that with her subjective probability of that outcome given that that action is performed. She should then sum these for all available actions and do the action with the greatest sum. How, Feldman asks, can she know how many alternative actions there are? How can she know what the outcomes are of those actions? How can she assign value to those outcomes? This is surely a troublesome objection. Although I think Feldman is wrong in thinking that it actually is *harder* to calculate expected value than it is to find out which act maximizes actual value, I believe that to be somewhat beside the point here.²⁴ As a defender of DTC I think there is no other option than to acknowledge that there will be cases where we cannot do the requisite mathematics. Feldman, however, concedes that there are cases where we can do it. Cases, for instance, with doctors such as Jill where we have a lot of empirical material at our disposal. When we have established the likelihood that a certain drug will have so and so effects we can use that information to assign values etc. There are however a huge number of cases where this is not the case, for instance, because they range over new

²³ 'Is Moral Obligation Objective or Subjective' (forthcoming). Cf. Gruzalski (1981: 167), Smart and Williams (1973: 86f.).

²⁴ Here is Feldman's idea: "The calculation of regular utility is problematic. To perform the calculation, one needs to know his alternatives, their total consequences until the end of time, and the values of those consequences. A daunting epistemic task! But the calculation of expected utility is even more problematic. To perform this calculation, one needs to know his alternatives, and for each alternative, one also needs to know every one of its possible outcomes, the actual value of each of those outcomes, and the probability of each outcome, given the act, on one's current evidence. One also needs to know some mathematical facts: the product, for each outcome, of its value and its probability; and the sum of those products. If the epistemic task in the actual utility case was daunting, this task is double-daunting." I simply cannot understand why the epistemic task is 'double-daunting'.

and partly unexplored alternatives. Is this *particularly* troubling for DTC or is it a situation where DTC and OC fare equally badly? I would admit as much as this: Feldman's objection somewhat takes the bite out of DTC. The initial attraction we had for the feature that we must focus on the ought most immediately relevant to action and the assumption that we must make use of the information available at the time of action, appears less attractive in the face of it. But if anything, Feldman's objection is a problem *about* consequentialism and even though he considers it to be a sufficient reason to reject DTC, I believe it is a structural problem that will apply to all forms of consequentialism. Whether it is an objection that ultimately will lead us to reject consequentialism (in any form) goes beyond the questions raised in this paper.

5.

I have argued that we should reject OC and embrace DTC. In doing this I have claimed that if we take seriously the idea that a normative theory should tell us what we should do and at the same time be action-guiding, we must adopt the idea that we should maximize expected value. An agent oughts_{subjective} to do the action that is rights_{subjective}, as I called it. The 'ought' that we found relevant for overall moral obligation was claimed not to be ambiguous. In fact that turned out to be the notion of oughts_{subjective}. DTC is still falling short of being a comprehensive normative theory since we have not given a plausible account of the 'epistemic ought'. That is surely an interesting task for the defenders of DTC. But save for this one feature, I still think it is the best version of consequentialism that we have at the moment.

References

- Bales, R. E. (1971) 'Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?', *American Philosophical Quarterly*, Vol. 8: 257-265
- Bergström, Lars. (1996) 'Reflections on Consequentialism', *Theoria* vol. LXII: 74-94
- Brandt, Richard. (1959) *Ethical Theory*, Englewood Cliffs: Prentice-Hall
- Brink, David. (1986) 'Utilitarianism and the Personal Point of View', *The Journal of Philosophy*, vol. 83: 417-438
- Carlson, Erik. (1995) *Consequentialism Reconsidered*, Dordrecht: Kluwer Academic Publishers
- Dancy, Jonathan. (2000) *Practical Reality*, Oxford: Oxford University Press
- Ewing, A. C. (1947) *The Definition of Good*, New York: The Macmillan Company
- Feldman, Fred. (Forthcoming) 'Actual Utility, the Objection from Impracticability, and the Move to Expected Utility'
- Gruzalski, Bart. (1981) 'Foreseeable Consequence Utilitarianism', *Australasian journal of Philosophy*, vol. 59: 163-176
- Harman, Gilbert. (1977) *The Nature of Morality*, New York: Oxford University Press

- Jackson, Frank. (1991) 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics*, vol. 101: 461-482
- Jackson, Frank. (1997) 'Which Effects?' in Jonathan Dancy (ed.), *Reading Parfit*, Oxford: Blackwell
- Mill, John Stuart. (1998) *Utilitarianism*, Roger Crisp (ed.), Oxford: Oxford University Press
- Moore, G. E. (1993 [1903]) *Principia Ethica*, Thomas Baldwin (ed.), rev. ed., Cambridge: Cambridge University Press
- Moore, G. E. (1965 [1912]) *Ethics*, Oxford: Oxford University Press
- Parfit, Derek. (1984) *Reasons and Persons*, Oxford: Clarendon Press
- Railton, Peter. (1984) 'Alienation, Consequentialism, and the Demands of Morality', *Philosophy and Public Affairs*, vol. 13: 134-171
- Sidgwick, Henry. (1981 [1907]) *The Methods of Ethics*, Indianapolis: Hackett
- Smart, J. J. C. (1973) 'An Outline of a System of Utilitarian Ethics' in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against*, Cambridge: Cambridge University Press
- Smith, Holly M. (1988) 'Making Moral Decisions', *Noûs*, Vol. 22: 89-108
- Sosa, David. (1993) 'Consequences of Consequentialism', *Mind* Vol. 102: 101-122
- Timmons, Mark. (2002) *Moral Theory: An Introduction*, Lanham: Rowman & Littlefield Publishers
- Zimmerman, Michael. (forthcoming) 'Is Moral Obligation Objective or Subjective?'
-