

Consequentializing Moral Theories*

DOUGLAS W. PORTMORE

BROADLY CONSTRUED, maximizing act-consequentialism (hereafter simply 'consequentialism') is the view that all acts are permissible or impermissible solely in virtue of how their outcomes rank relative to those of the available alternatives on some transitive, evaluative ordering of outcomes such that an act is permissible if and only if no available alternative outcome ranks higher than its outcome on that ordering.¹ An act's outcome is also to be construed broadly so as to include everything that would be the case were the act to be performed.² When 'consequentialism' and 'outcome' are both construed broadly, it turns out, as I will later argue, that any remotely plausible nonconsequentialist theory can be consequentialized. The recipe for consequentializing a nonconsequentialist theory is simple: Take whatever considerations that the nonconsequentialist theory holds to be relevant to determining the deontic status of an action and insist that those considerations are relevant to determining the proper ranking of outcomes. In this way, the consequentialist can produce an ordering of outcomes that when combined with her criterion of rightness yields the same set of deontic verdicts that the nonconsequentialist theory yields—that is, for any deontic predicate ('permissible', 'impermissible', 'obligatory', 'supererogatory', etc.), the

* This is the penultimate draft (4/22/06) of a paper that is forthcoming in the *Pacific Philosophical Quarterly*.

¹ Regarding terminology, let me say, first, that although I'll be using the terms 'consequentialism' and 'consequentializing', I might have just as well used the terms 'teleology' and 'teleologizing'. Some authors who discuss the topic of this paper prefer to use the term 'teleology' as opposed to 'consequentialism' (e.g., Broome 1991, 4-6), whereas others have the opposite preference (e.g., Brown 2004, Dreier 1993, and Louise 2004). My decision to go with 'consequentialism' and 'consequentializing' is mostly arbitrary, and nothing I say in this paper hangs on my choice of terms. Second, what I and others (e.g., Brown 2004 and Dreier 1993) call 'consequentializing' is sometimes called 'the consequentialist umbrella' (Louise 2004) or 'the consequentialist vacuum cleaner' (McNaughton and Rawling 1991; Ridge forthcoming).

² On why consequentialists should construe an act's outcome to include not merely its causal consequences but everything that would be the case were it to be performed, see Broome 1991, 3-4; Scheffler 1982, 1-2; Sosa 1993, 101-2; and Williams 1973, 86-87.

resulting consequentialist counterpart theory and the original nonconsequentialist theory will be in perfect agreement as to the set of actions that are in the extension of that predicate (Brown 2004, 2).³

To illustrate how consequentializing works, suppose that a nonconsequentialist holds that an agent's intentions are relevant to determining whether or not her actions are morally permissible. More specifically, suppose that this nonconsequentialist holds that although it is, say, permissible to kill one as a merely foreseen but unintended consequence of doing what will save five others, it is impermissible to intend to kill one for the sake of saving five others. In this case, all the consequentialist needs to do to yield these same two deontic verdicts is to hold that although the outcome in which an agent has killed someone as a merely foreseen but unintended consequence of saving five others ranks higher than the outcome in which the agent has instead allowed those five others to die, the outcome in which an agent has intentionally brought about someone's death in order to save five others ranks lower than the outcome in which the agent has instead allowed those five others to die. Of course, a lot more will need to be said before we can conclude that the consequentialist can employ this consequentializing strategy to accommodate all the various sorts of considerations that nonconsequentialists take to be relevant to determining an act's deontic status, but that will come later. Let me begin instead by explaining the structure of the paper.

In section 1, I explain the two aims that together constitute the motivation for consequentializing and rebut Mark Schroeder's (forthcoming) recent arguments that the consequentializing project cannot succeed since, as he argues, these two aims cannot be simultaneously achieved. I argue that what follows from Schroeder's arguments is not that

³ This doesn't entail that the original nonconsequentialist theory and the resulting consequentialist counterpart theory will agree about *everything*. They will agree on which considerations (or factors) are deontically relevant, but they may, nevertheless, disagree about the rationale (or the foundation) for why these and not others are the relevant factors (Kagan 1992). For instance, a contractualist and a consequentialist might both agree that the production of aggregate utility (i.e., aggregate welfare) is the only consideration relevant to determining an act's deontic status, but whereas a consequentialist will accept this only if she believes that aggregate utility is the only thing relevant to the ranking of outcomes, a contractualist will accept this only if she believes that the principle of utility is the only principle "that no one could reasonably reject as a basis for informed, unforced general agreement" (Scanlon 1998, 153). Thus, even if they agree on what the morally relevant factors are, they will, nevertheless, disagree at the foundational level. At any rate, I'll have more to say about this in section 5.

the consequentializing project is doomed to fail, but that we must rethink how consequentialism is to be defined as well as rethink what it is about act-utilitarianism that seems so compelling. In sections 2 and 3 respectively, I take up these two tasks. In section 4, I argue that, for any remotely plausible nonconsequentialist theory, there is a consequentialist counterpart theory that is deontically equivalent to it—that is, the two theories will be extensionally equivalent with respect to their deontic verdicts regarding act-tokens.⁴ I call this the ‘Deontic Equivalence Thesis’ or ‘DET’ for short. I go beyond previous attempts to demonstrate DET by showing that, for any remotely plausible nonconsequentialist theory, we can construct a consequentialist counterpart theory that not only mimics its verdicts about what’s permissible but also mimics its verdicts about what’s supererogatory. What’s more, I show that consequentialism can accommodate agent-centered options. In section 5, I take issue with two conclusions that other philosophers have drawn from DET: *(i)* that we are all consequentialists (Dreier 1993 and Louise 2004) and *(ii)* that consequentialism is empty (Brown 2004). I argue that neither is the case. Lastly, in section 6, I argue that although the consequentializer will need to appeal to our considered moral convictions in determining how to rank outcomes, this in no way renders the resulting consequentialist position circular or uninformative. The net effect of these arguments is, I hope, to demonstrate that the consequentializing project is a viable, and indeed promising, one.

1. Why consequentialize?

First and foremost, consequentializing is a strategy for defending consequentialism against the sorts of objections that nonconsequentialists have typically leveled against traditional forms of consequentialism, such as act-utilitarianism. Nonconsequentialists argue that act-utilitarianism fails as a criterion of rightness because there are other factors besides aggregate utility (i.e., aggregate welfare) that are relevant to determining an act’s deontic status. The consequentializing strategy enables the consequentialist to meet this objection: All the considerations that the nonconsequentialist cites as being relevant to determining an act’s deontic

⁴ Throughout this paper, whenever I refer to deontic verdicts, I’m referring to verdicts regarding act-tokens—concrete, individual acts that are performable by a particular agent at a particular time in a particular set of circumstances. Thus, as we’ll see in section 5, theories that are extensionally equivalent in their deontic verdicts can differ in their fundamental moral principles and, thus, in their rationales (i.e., their foundations) for those verdicts.

status are, the consequentialist argues, considerations relevant to determining how outcomes are to be ranked and, consequently, relevant to determining an act's deontic status even on consequentialism. But the motivation for consequentializing is not merely to avoid counter-intuitive implications. After all, many nonconsequentialist theories avoid counter-intuitive implications. The point of consequentializing is to avoid counter-intuitive implications while remaining consequentialist, for there is supposedly something at the heart of a theory like act-utilitarianism that is deeply compelling, and it is not its overly narrow welfarist conception of the good but rather its consequentialism. Indeed, many nonconsequentialists acknowledged that there is something deeply compelling about consequentialism and that this explains the "spellbinding force" that act-utilitarianism has had over even those who refuse to accept it (Foot 1985, 196-198). So the motivation for consequentializing is to keep what's compelling about act-utilitarianism (i.e., its consequentialism) while avoiding what's problematic about it (i.e., its abundant counter-intuitive implications). The consequentializing project is, then, to come up with a theory that achieves these two aims by combining consequentialism's criterion of rightness with a more sophisticated account of how outcomes are to be ranked (an account that is, at least, more sophisticated than the account that welfarism provides).

Mark Schroeder, in a series of recent papers (forthcoming, 2005a, 2005b), argues that the consequentializing project cannot succeed. Following others (including my past self), Schroeder identifies what's so compelling about a theory like act-utilitarianism as the idea that it is always permissible to bring about the most good. He calls this the "Compelling Idea." And Schroeder claims, again following others, that in order to eschew the counter-intuitive implications that plague act-utilitarianism, the consequentializer will need to accommodate agent-centered constraints.⁵ Schroeder then argues that in order to achieve the goal of accepting the Compelling Idea, the consequentializer must define rightness in terms of the ordinary language word 'good' (the monadic predicate), but that in order to achieve the goal of accommodating agent-centered constraints, the

⁵ These impose limits on an agent's freedom to permissibly pursue the best consequences, impersonally construed, and they include both restrictions and special obligations. Restrictions prohibit the performance of certain act-types (e.g., murder) even in order to prevent numerous others from each performing that very same act-type. Special obligations arise due to past promises or through institutionally defined roles, and it is wrong to violate such obligations even if doing so will produce the best consequences, impersonally construed. See Kagan 1989, esp. 4.

consequentializer must define rightness in terms of the dyadic ‘good-relative-to’ relation—that is, the consequentializer must adopt agent-relative consequentialism, the view that, for all x , x ought always to do what will bring about the most good-relative-to x .⁶ The problem is, as Schroeder sees it, there is no viable story about how the good-relative-to relation and the ordinary English word ‘good’ get connected such that it will be true to say that the agent-relative consequentialist (which is the only sort of consequentialist that can accommodate agent-centered constraints) accepts the Compelling Idea. That is, Schroeder argues that there is no plausible semantics of the ordinary language word ‘good’ (the monadic predicate) that renders agent-relative consequentialism compatible with the Compelling Idea. Therefore, there is, according to Schroeder, no way for the consequentializer to accommodate both the Compelling Idea and agent-centered constraints.

For the sake of argument, let us grant Schroeder’s claim that there is no viable story about how the ordinary English word ‘good’ and the good-relative-to relation get tied together such that it will be true to say that the agent-relative consequentialist accepts the Compelling Idea. So although the agent-relative consequentialist holds that it is always permissible for an agent to bring about the most good-relative-to her, this is not the Compelling Idea. The Compelling Idea is that it is always permissible for an agent to bring about the most good, where ‘good’ here is the ordinary monadic predicate of English, not the dyadic good-relative-to relation. Nevertheless, I will argue that what follows from Schroeder’s arguments is not that the consequentializing project is doomed to fail, but that philosophers, myself included, have misidentified what it is that is so compelling about act-utilitarianism. Here, it’s important to note that the compelling idea that lies at the heart of act-utilitarianism is supposed to lie also at the heart of other consequentialist theories, broadly construed, such as ethical egoism. But as Schroeder admits, ethical egoism cannot accommodate the Compelling Idea anymore than agent-relative consequentialism can. He (forthcoming) says,

Ethical egoism is a teleological view [read: consequentialist view] that employs the *good for* concept *instead* of the one expressed by ‘good’ when used as a monadic predicate. It says that rather than

⁶ The same idea is sometimes put as follows: for all x , x ought always to do what will bring about the most good _{x} , where to judge that p ’s being the case is good _{x} is to judge that x would desire that p if x had a psychology that eludes all forms of rational criticism—see Smith 2003. The term ‘good _{x} ’ is just a variant of ‘good-relative-to x ’.

doing what will bring about the most good, you should do what will bring about the most of what is good *for you*. ...But since *good* and *good for* express different concepts, the idea that egoists find compelling turns out not to be at all the same idea that consequentialists [read: agent-neutral consequentialists] find Compelling.⁷

So if Schroeder is right, then it was a mistake to identify the compelling idea as the idea that it is always permissible for an agent to bring about the most good, for the compelling idea is supposed to be something that act-utilitarianism and ethical egoism have in common. Furthermore, if Schroeder is right, we can't define consequentialism in terms of the ordinary English word 'good', not if we want, as consequentializers do, 'consequentialism' to be an umbrella term that includes both act-utilitarianism and ethical egoism. In the next section, then, I'll explain how we might properly define 'consequentialism' in light of Schroeder's arguments, and, in the subsequent section, I'll identify what it is that is so compelling about *all* consequentialist theories, including both act-utilitarianism and ethical egoism.

2. Defining Consequentialism

Ethical egoism, act-utilitarianism, and what I've called "agent-relative consequentialism" are all consequentialist theories—at least, they are if we are to use the term 'consequentialism' in the same sense that Brown, Dreier, and Louise have claimed that all plausible moral theories can be consequentialized. But if this is the case, then we cannot define 'consequentialism' in terms of the concept expressed by 'good' when used as a monadic predicate, for, as Schroeder has argued, these three theories employ three distinct evaluative concepts. Whereas act-utilitarianism makes the deontic status of an action a function of what's *good*, ethical egoism makes the deontic status of an action a function of what's *good for* the agent. Different still, agent-relative consequentialism makes the deontic status of action a function of what's *good-relative-to* the agent. So we can't define 'consequentialism' as the theory that holds that an act is permissible just in case it produces the most good. For although that's what act-utilitarians hold, ethical egoists, by contrast, hold that an act is permissible

⁷ Following my 2005b, Schroeder uses a different set of terms than I do here, but the terms are easily translatable. Schroeder's 'teleology' is what I now call 'consequentialism'. What Schroeder calls 'consequentialism' is what I now call 'agent-neutral consequentialism'.

just in case it produces the most of what's good for the agent. Different still, the agent-relative consequentialist holds that an act is permissible just in case it produces the most of what's good-relative-to the agent.

It is easy to see that these three accounts of permissibility don't amount to the same thing. To say that O. J. Simpson's acquittal was what was *best for him* is not to say that it was what was *best*. And to say that the state of affairs, S1, where I, A1, kill one innocent person, P1, to prevent five other agents, A2-A6, from each killing an innocent person, P2-P6, is *worse-relative-to me* than the state of affairs, S2, where I refrain from killing P1 and thereby allow A2-A6 to kill P2-P6 is not to say that S1 is *worse for me*. It's possible, for instance, that P1 is someone I detest, whereas P2-P6 are all people whom I love; in that case, S2 might be *worse for me*. And to say that S1 is *worse-relative-to me* is not to say that S1 is *worse*, for it may be worse that more innocent people are murdered. So if Schroeder is right, we must define 'consequentialism' in a way that's neutral with respect to these three different evaluative concepts. In what follows, I suggest that the best way to do this is to define 'consequentialism' in terms of an ordering of outcomes with respect to what the agent has reason to prefer.

Let 'P' be a variable standing for a person who must, at t , perform one out of a set of mutually exclusive but jointly exhaustive acts $a_1 \dots a_n$ with corresponding outcomes $o_1 \dots o_n$.⁸ A theory is consequentialist if and only if there exists some transitive ordering of outcomes, the $\succ_{P/t}$ ordering, such that the following two conditions are met:

- (1) The Determination Condition: whether an act is permissible or impermissible is fully and ultimately determined by how its outcome ranks relative to those of the available alternatives on the $\succ_{P/t}$ ordering such that an act is permissible if and only if no available alternative outcome ranks higher than its outcome on that ordering, and
- (2) The Preferability Condition: $o_i \succ_{P/t} o_j$ if and only if facts about o_i and o_j make it fitting for P, at t , to prefer o_i to o_j —that is, if and only if P, at t , has better object-given reasons to prefer o_i to o_j than to prefer o_j to o_i .⁹

⁸ Of course, since P must, at t , perform one of $a_1 \dots a_n$, this set must include the act of not performing any voluntary bodily movement at t .

⁹ ' $o_i \succ_{P/t} o_j$ ' means ' o_i ranks higher than o_j on the $\succ_{P/t}$ ordering'. Some might wonder why I define consequentialism in terms of the \succ relation rather than the \succeq relation. The reason is that certain consequentialist theories, such as Schefflerian Utilitarianism (described in

Both of these require some explanation. So let me take each in turn. First, given the determination condition, the consequentialist is committed to the view that all acts are right or wrong *in virtue* of how their outcomes rank relative to those of the available alternatives on the $\succ_{v/t}$ ordering. Moral theories are supposed to tell us what makes acts right or wrong, and the determination condition ensures that consequentialism fulfills this function. Also, the determination condition establishes that, on consequentialism, the ranking of outcomes is prior to the determination of an act's deontic status in that one must first determine how a set of alternative outcomes are ranked before one can determine whether the acts that produce them are permissible or not.¹⁰ Thus, the consequentialist must deny that whether $o_i \succ_{v/t} o_j$ is determined by whether P 's doing a_i at t is deontically superior to P 's doing a_j at t . (Let me just stipulate that one act is deontically superior to another if and only if the former is permissible and the latter is impermissible.) A nonconsequentialist, by contrast, can hold that how a set of alternative outcomes are to be ranked is just determined by the deontic statuses of the acts that produce them.

Second, the preferability condition is essential to a definition of consequentialism, because it's not enough that there be some ordering such that the deontic status of an action is a maximizing function of that ordering. That's true of almost any moral theory.¹¹ What makes a theory distinctively consequentialist, then, is that the relevant ordering is an evaluative one, one such that the agent has better *object-given reasons* to

section 4 below), cannot be represented by a transitive \succeq relation. Of course, an ordering stated in terms of the \succ relation won't be able to distinguish instances where outcomes are equal from instances where outcomes are incommensurable. Nevertheless, it seems that the consequentialist needn't make such a distinction, for consequentialism holds that an act is permissible if and only if no available alternative outcome ranks higher than its outcome. Thus, whether an alternative outcome, o_j , is equal or incommensurable with another, o_i , makes no difference, for in neither case will the availability of o_j make a_i impermissible.

¹⁰ Note that "determination has an asymmetry that derives from an implied in-virtue-of or explanation relation. That is, if F determines G , it is asymmetrically in virtue of, hence explained by, having F that something has G " (Post 1999).

¹¹ As Robert Nozick (1968; 1974, 28-39) has pointed out, almost any moral theory can be defined as a maximizing theory, a theory that holds that there is some transitive ordering such that an action is permissible just in case it ranks highest on that ordering. Take, for instance, the divine command theory. "Define a function, f , taking actions as arguments, such that $f(ac) = 1$ just in case ac satisfies God's commands, and $f(ac) = 0$ otherwise. The divine command theory is representable as a maximizing theory, since it judges an action permissible just in case it has a maximal f value" (Vallentyne 1988b, 253-4).

prefer the higher-ranked outcome to the lower-ranked outcome.¹² P's reason to prefer o_i to o_j is an object-given reason (as opposed to a state-given reason) if and only if it is provided by facts about o_i and o_j (the objects of P's preference) as opposed to facts about P's state of preferring o_i to o_j .¹³ So, for instance, whereas the fact that o_i contains more aggregate utility than o_j provides P with an object-given reason to prefer o_i to o_j , the fact that an evil demon will destroy the world and everyone in it unless P prefers o_i to o_j provides P with only a state-given reason to prefer o_i to o_j . According to the preferability condition, it is only P's having better object-given reasons to prefer o_i to o_j that implies that it is fitting for P to prefer o_i to o_j .¹⁴ By contrast, P's having better state-given reasons to prefer o_i to o_j implies only that it is fitting for P to prefer being in the state of preferring o_i to being in the state of preferring o_j . To sum up, it is not enough that P has better state-given reasons to prefer o_i to o_j . For it to be fitting for P to prefer

¹² According to a very influential tradition, evaluative concepts (e.g., desirable, admirable, disgusting, shameful, etc.) are to be analyzed in terms of certain kinds of reasons for adopting certain kinds of pro- or con-attitudes (e.g., desire, admiration, disgust, shame, etc.) toward the objects of which the evaluative concepts are being predicated. Thus, on this sort of analysis, X is E (where E stands for some evaluative concept) if and only if there are facts about X that constitute reasons of a certain sort for adopting a certain pro- or con-attitude toward X. This kind of approach to analyzing evaluative concepts goes by many different names—including “the buck-passing account,” “the fitting-attitudes analysis,” “dispositionalism,” and “sentimentalism”—and has been advocated in some form or another by a number of different philosophers, including Franz Brentano (1989), A. C. Ewing (1947), Allan Gibbard (1990), John McDowell (1985), Derek Parfit (2001), Thomas Scanlon (1998) and David Wiggins (1987). I don't wish to take a stand here on whether this is the correct way to analyze such concepts, for I don't intend to give an analysis of any evaluative concept here. Rather, I merely want to point out that my way of ensuring that the consequentialist's ordering of outcomes is an evaluative one dovetails quite nicely with this sort of analysis.

¹³ The distinction between object-given reasons and state-given reasons comes from Parfit 2001. See also Parfit forthcoming.

¹⁴ Rabinowicz and Rønnow-Rasmussen argue that it won't suffice to say that the relevant reasons are object-given reasons as opposed to state-given reasons, for they argue that, for any fact about the state of preferring o_i , there is always a corresponding fact about o_i —see their 2004, 404-408. So, for instance, if a fact about my preferring o_i is that it will shield me from punishment by an evil demon, then it is a fact about o_i itself that my preferring it will shield me from punishment by an evil demon. Thus they claim that, for any state-given reason, there is always a corresponding object-given reason. But see Stratton-Lake 2005 for a decisive reply. Stratton-Lake argues that even if there is a corresponding fact about o_i for every fact about my preferring o_i , such facts about o_i fail to provide reasons. Thus, it's false that, for any state-given reason, there is always a corresponding object-given reason. In any case, if the reader prefers, he or she may take my ‘object-given reasons’ to be a placeholder for whatever the right kind of reasons are.

o_i to o_j , P must have better object-given reasons to prefer o_i to o_j . And let me clarify why I use the phrase “better reason” as opposed to, say, “more reason”: As I see it, the former allows for a greater range of possibilities, such that P might have better reason to prefer o_i , not only because P’s reasons for preferring o_i are weightier than P’s reasons for preferring o_j , but also because P’s reasons for preferring o_j are trumped, silenced, undermined, excluded, or bracketed off by other reasons.¹⁵

It is important to note that the above definition is compatible with a consequentialist holding that what an agent ought to do and what an agent ought to bring it about that she prefers can come apart in just the way that some consequentialists have alleged when defending consequentialism against objections concerning integrity, friendship, and the like. For instance, a utilitarian might hold that although an agent ought not to act so as to benefit a friend to a lesser extent when she could instead benefit a stranger to a greater extent, she ought to bring it about (if she can) that she prefers an outcome in which a friend is benefited to a lesser extent to an outcome in which a stranger is benefited to a greater extent. This is because a consequentialist can be a consequentialist not only about acts, but also about desires, beliefs, motives, preferences, dispositions, etc. In each case, the consequentialist can hold that an agent ought to bring it about that ϕ if and only if no available alternative outcome ranks higher on the \succ_{vt} ordering than the outcome that would result from her bringing it about that ϕ , where ‘ ϕ ’ is a placeholder for anything that the agent can bring about: that she has performed a given act, that she has a certain belief, that she has a certain preference, etc. So although the above definition of consequentialism doesn’t allow that what an agent ought to do can come apart from an agent has best object-given reasons to prefer, it does allow that what an agent ought to do can come apart from what she ought to bring it about that she prefers, for a consequentialist about what preferences an agent ought to bring it about that she has will necessarily appeal to state-given reasons for having those preferences, i.e., to reasons that stem from the value of the outcomes that will result from her bringing it about that she is in those preference-states.

Note also that I don’t intend for the above definition to capture the sense in which all, or even most, philosophers use the term ‘consequentialism’. I take ‘consequentialism’ to be a term of art and thus open to being stipulatively defined, as I’ve done here. In any case, it seems that

¹⁵ For a discussion of excluding, see Raz 1975, pp. 37-39. For a discussion of bracketing off, see Scanlon 1998, pp. 50-54.

‘consequentialism’ must be defined along these lines if it is going to encompass all the various theories that consequentializers such as Brown, Dreier, and Louise include under the rubric, theories such as ethical egoism, act-utilitarianism, and agent-relative consequentialism. And it’s important to include all such theories, for, as we’ll see in the next section, they all share the same compelling idea that lies at the heart of act-utilitarianism. I’m in agreement, then, with Walter Sinnott-Armstrong (2003) when he says, “the term ‘consequentialism’ seems to be used as a family resemblance term to refer to any descendent of classic utilitarianism that remains close enough to its ancestor in the important respects.” And, to my mind, sharing this compelling idea is *the* most important respect.

But before moving on to discuss what this compelling idea is, it is important to note that, on my proposed definition, what distinguishes one version of consequentialism from another is how they each specify the $\succ_{P/t}$ ordering. To illustrate, consider the following four versions of consequentialism:

- EP** Egoism of the Present: $o_i \succ_{P/t} o_j \leftrightarrow U_{P/t}(o_i) > U_{P/t}(o_j)$, where $U_{P/t}(o_i)$ denotes what P ’s momentary utility would be at t were o_i to obtain.¹⁶
- EE** Ethical Egoism: $o_i \succ_{P/t} o_j \leftrightarrow U_P(o_i) > U_P(o_j)$, where $U_P(o_i)$ denotes what P ’s lifetime utility would be were o_i to obtain.
- AU** Act-Utilitarianism: $o_i \succ_{P/t} o_j \leftrightarrow U(o_i) > U(o_j)$, where $U(o_i)$ denotes what the total aggregate utility (across time and individuals) would be were o_i to obtain.
- SOU** Self/Other Utilitarianism: $o_i \succ_{P/t} o_j \leftrightarrow$ both (1) $U(o_i) > U(o_j)$ and (2) $U_{-P}(o_i) > U_{-P}(o_j)$, where $U_{-P}(o_i) = U(o_i) - U_P(o_i)$.¹⁷

These four theories demonstrate just how diverse and wide-ranging consequentialism can be. We see that a consequentialist theory can appeal to various different evaluative concepts (e.g., good, good for, or good-relative-to) and can be either agent-neutral or agent-relative and either temporally-neutral or temporally-relative. A consequentialist theory is agent-relative if and only if the right-hand side of the bi-conditional must

¹⁶ This theory is discussed by Derek Parfit—see his 1984, 134.

¹⁷ See Ted Sider’s 1993.

unavoidably make reference to P , and a consequentialist theory is temporally-relative if and only if the right-hand side of the bi-conditional must unavoidably make reference to t . Thus EP is both agent-relative and temporally-relative, whereas EE is agent-relative but temporally-neutral. Both theories, though, appeal to the same evaluative concept in holding the deontic status of an action to be a function of what's *good for* the agent. By contrast, AU is both agent-neutral and temporally-neutral and makes the deontic status of an action a function of what's *good*. SOU is agent-relative but temporally-neutral and makes the deontic status of an action a function of what's *good-relative-to* the agent.¹⁸ So the great merit of this way of defining consequentialism is that it is able to encompass the vast variety of particular theories that those interested in the consequentializing project have considered to be consequentialist. Having defined consequentialism in this way, the success of the consequentializing project now hangs on whether it is correct to diagnose what it is about act-utilitarianism that is so compelling as something that all consequentialist theories, so defined, have in common, an issue to which I now turn.

3. What's so compelling about act-utilitarianism?

What is it about act-utilitarianism that philosophers find so compelling and even "spellbinding"? In other words, what explains: (1) why act-utilitarianism has persevered despite its implications being so wildly at odds with our most firmly held moral convictions;¹⁹ (2) why it tends "to haunt even those of us who will not believe in it" (Foot 1985, 196), and (3) why "the move to rule utilitarianism seems to be an unsatisfactory answer to the problem of reconciling utilitarianism with common moral opinion" (Foot 1985, 198)? Those interested in the consequentializing project had better hope that the explanation doesn't lie with act-utilitarianism's endorsement of the idea that it is always permissible to produce the most *good*, for, as Schroeder has argued, the consequentializing project won't get

¹⁸ SOU directs agents to produce outcomes for which there is no available alternative outcome that contains both more aggregate utility and more aggregate utility for others (i.e., for everyone other than x). This entails neither maximizing what's *good* nor maximizing what's *good for x*. So if SOU is a version of consequentialism (i.e., a version of maximizing act-consequentialism) at all, it's a version of agent-relative consequentialism, directing agents to maximize what's *good-relative-to x*, where x 's producing the most of what is *good-relative-to x* entails x 's producing an outcome for which there is no available alternative that contains both more aggregate utility and more aggregate utility for others.

¹⁹ Speaking of act-utilitarianism in 1973, Bernard Williams predicted that the "day cannot be too far off in which we hear no more about it" (1973, 150).

very far if that's the case. In order to accommodate agent-centered constraints, the consequentializer must deny that agents ought always to produce the most *good* and accept instead that each agent ought always to produce the most of what's *good-relative-to* her. Fortunately for the consequentializing project, though, the compelling idea behind act-utilitarianism is supposedly something that it shares in common with ethical egoism and that can't be, if Schroeder is right, that it is always permissible to bring about the most *good*, for, as we have seen, ethical egoism denies this. What, then, is the compelling idea that both act-utilitarianism and ethical egoism have in common? It is, I think, what Scheffler calls "the maximizing conception of rationality." As Scheffler puts it,

The kind of rationality that consequentialism seems so clearly to embody, and which makes so much trouble for views that incorporate agent-centered restrictions, is what we may call *maximizing* rationality.^[20] The core of this conception of rationality is the idea that if one accepts the desirability of a certain goal being achieved, and if one has a choice between two options, one of which is certain to accomplish the goal better than the other, then it is, *ceteris paribus*, rational to choose the former over the latter. (1985, 414)

Now, according to Scheffler, deontologists are not only committed to each agent having the agent-relative goal of ensuring that she not violate a restriction, but also committed to each agent having the agent-neutral goal of minimizing violations of the restrictions. Given this agent-neutral goal, theories that are committed to agent-centered restrictions can seem paradoxical in that they sometimes prohibit an agent from violating a restriction even when doing so will better achieve this goal, even when doing so will minimize the number of violations of that restriction overall.²¹

²⁰ As Scheffler defines 'consequentialism', it is necessarily agent-neutral and, therefore, incapable of accommodating agent-centered restrictions. Had he noted the possibility of agent-relative consequentialism, he might have seen that it could incorporate agent-centered restrictions while at the same time embodying the maximizing conception of rationality.

²¹ Of course, as Scheffler admits, "it can sometimes be rational to act in such a way as to worse achieve one goal if that will make it possible to better achieve another. Since that is so, it might be said, views that include agent-centered restrictions need not come into conflict with maximizing rationality when they tell us to further the agent-relative goal of not violating the restrictions ourselves at the expense of the non-relative goal of minimizing

Interestingly, though, Scheffler doesn't view ethical egoism as being paradoxical. He says,

Notice that egoism, by contrast, *does* seem committed exclusively to agent-relative goals. It assigns each person the agent-relative goal of maximizing his own advantage. And since it does not purport to assess action from a point of view which is concerned with more than just the interests of the agent, it is not committed in the way deontology is to presenting as desirable any non-relative goal whose maximum accomplishment it then prohibits. That is why *it does not for a moment seem paradoxical for the egoist to say that one ought to maximize one's own advantage even if that means that fewer people overall will be able to maximize theirs.* (1985, 416)

Ethical egoism seems, then, to endorse the maximizing conception of rationality. To see this more clearly, it might be helpful to consider what an egoistic theory that didn't accept the maximizing conception of rationality would look like. Consider, then, two nonconsequentialist versions of egoism: deontological egoism and rule egoism. Both accept that it is always fitting for an agent to prefer an outcome in which she has more utility to one in which she has less utility, and yet both hold that it is sometimes wrong for an agent to act so as to bring about the outcome in which she has maximal utility. For instance, on deontological egoism, there are certain types of acts, say, self-sacrificing acts, that are intrinsically immoral such that it would be wrong for an agent to commit one self-sacrificing act now even to prevent herself from committing more numerous and equally self-sacrificing acts in the future. This view seems paradoxical, for if the sole goal is to maximize one's utility and thus, derivatively, also to minimize one's self-sacrifices, then why insist that it would be wrong to maximize one's advantage by performing a single self-sacrificing act now, thereby minimizing the total number of equally self-sacrificing acts one will perform overall?

violations of the restrictions" (1985, 417). But Scheffler argues that this latter claim is not persuasive, for, as he contentiously assumes, the agent-relative goal of not violating the restrictions ourselves is "derivative from, and given life by," the non-relative goal of minimizing violations of the restrictions (1985, 417). In any case, the point here is not to agree with Scheffler's contention that views that incorporate agent-centered restrictions are paradoxical but rather to argue that he has correctly identified what it is that is so compelling about act-utilitarianism as the embodiment of the maximization conception of rationality.

On rule egoism, an agent ought always to follow the set of rules that, if internalized by her, would lead her to produce more utility for herself than any other alternative set of rules would (Kagan 1992, 238). This view also seems paradoxical. To illustrate, suppose that the ideal set of rules includes a rule that prohibits donating one's money to charity. On this view, then, an agent ought not to donate her money to charity even if she knows that, in this instance, doing so will maximize her utility. But if the sole goal is to maximize one's utility, if it is always fitting for an agent to prefer more to less utility for herself, then it just seems paradoxical for such a theory to insist that it is wrong for her to donate her money to charity when doing so will clearly maximize her utility.²²

Unlike deontological egoism and unlike rule egoism, ethical egoism (the consequentialist theory) seems to share the same compelling feature that act-utilitarianism possesses: the maximizing conception of rationality. It is the attractiveness of the maximizing conception of rationality that explains why act-utilitarianism has had such a spellbinding force over contemporary moral philosophy and why the move to consequentialize has seemed more compelling than the move to rule-utilitarianism as a possible solution to act-utilitarianism's counter-intuitive implications. As William Shaw notes, "Consequentialism's goal-oriented, maximizing approach to ethics coheres with what we implicitly believe to be rational conduct in other contexts, in particular, when it comes to assessing prudential behavior" (2006,16). No one would seriously advocate a nonconsequentialist version of prudence, such as deontological prudence or rule prudence. As far as I am aware, no one holds that there are certain acts that are intrinsically imprudent such that it would be imprudent to perform such an act even when doing so is clearly what would best advance one's personal interests. Nor am I aware of anyone who holds that it is sometimes imprudent to perform an act that one knows would best advance one's personal interests just because it violates some ideal code of prudential rules. Why, then, would the moral sphere be the only sphere of

²² Note, though, that Brad Hooker has argued that rule consequentialism is not guilty of this sort of puzzling rule worship, for, according to Hooker, the rule consequentialist is *not* committed to the maximization of the good as an overarching goal. He says that the best argument for rule consequentialism is not that it derives from an overarching commitment to maximize the good, but that "it does a better job than its rivals of matching and tying together our moral convictions" (2000, 101) But rule egoism cannot make the same claim; it is not nearly as successful as its rivals at matching and tying together our moral convictions. Thus the best argument for rule egoism would have to be that it derives from an overarching commitment to maximize what's good for the agent. Consequently, rule egoism is guilty of rule worship.

rational conduct in which the maximizing conception of rationality didn't hold?

But what precisely is the maximizing conception of rationality? I think that we can improve on Scheffler's explication of it in at least two respects. For one, what Scheffler says is ambiguous. When at the end of the first of the above two quotes Scheffler says, "it is rational to choose the former over the latter," Scheffler could mean by 'rational' either 'not irrational' or 'rationally required'. If he means 'not irrational', then satisficing consequentialists can accept the maximizing conception of rationality (and I don't take this to be an objection), for in that case the adherent of the maximizing conception of rationality need only accept that it is never irrational (or impermissible) to produce the optimal (i.e., highest ranked) outcome.²³ And a satisficing consequentialist can hold that it is sometimes rational (or permissible) to produce a sub-optimal outcome while also acknowledging that it is never irrational (or impermissible) to produce the optimal outcome. Indeed, satisficing consequentialists typically hold that producing the optimal outcome is, although not required, supererogatory. Furthermore, it seems right to hold that satisficing consequentialism shares with maximizing act-utilitarianism the same compelling idea. After all, what's so puzzling about a view like rule-utilitarianism is that it holds both that it is always fitting for us to prefer that maximal aggregate utility is achieved but that sometimes we ought not to act so as to achieve this. By contrast, there is nothing nearly so puzzling, if puzzling at all, about satisficing act-utilitarianism, which holds only that we are not always required to act so as to achieve maximal aggregate utility. What's compelling about maximizing act-utilitarianism, then, is not that it holds that we are always required to do what will best achieve our theory-given goal(s), but that it holds that we are never prohibited from doing what will best achieve our theory-given goal(s). On this latter point, satisficers and maximizers agree.

Another way that we might do better than Scheffler in explicating what he calls "the maximizing conception of rationality" is to eschew talk of rationality, which is potentially misleading since what we're really interested in is permissibility. Consider that whether it is rational for an agent to choose what will worse accomplish a goal depends on whether

²³ As Mark Schroeder has pointed out to me, this might explain why, in contrast to agent-centered restrictions (restrictions on maximizing the good), Scheffler doesn't find agent-centered options (options to permissibly refrain from maximizing the good) at all puzzling.

she believes that it will worse accomplish that goal.²⁴ So an act-utilitarian can certainly hold that it is rational (i.e., not irrational) for an agent to choose what will worse accomplish maximal aggregate utility provided the agent believes that what she is doing will best accomplish that goal. What we're really interested in, then, is whether it is ever impermissible (in the 'knowledge-supposing sense'²⁵) to bring about an outcome that is in fact ranked higher than any available alternative. According to what I will call "the permissibility-of-maximizing view," the answer is "no."

PMV It is always permissible (in the knowledge-supposing sense) for an agent to act so as to bring about the highest ranked available outcome, i.e., the outcome that she has better object-given reasons to prefer above all other available alternatives.²⁶

PMV is preferable to other ways we might specify the compelling idea that lies at the heart of act-utilitarianism, such as either (a) it is always permissible to bring about the most good or (b) it is never permissible to

²⁴ I'm following Parfit here. He says, "While reasons are provided by the facts, the rationality of our desires and acts depends instead on our beliefs. When we know the relevant facts, these questions ['What do we have most reason to want, and do?' and 'What is it most rational for us to want, and do?'] have the same answers. But if we are ignorant, or have false beliefs, it can be rational to want, or do, what we have no reason to want, or do." (2001, 17)

²⁵ An act is permissible in the knowledge-supposing sense if and only if it is not one that a perfectly virtuous person with knowledge of all relevant facts would characteristically refrain from doing in the circumstances. An act is permissible in the evidence-relative sense if and only if it is not one that a perfectly virtuous person with only the evidence at hand would characteristically refrain from doing in the circumstances. Note that accepting one of these two bi-conditionals doesn't make one a virtue ethicist, for the virtue ethicist accepts the stronger claim that what *makes* an act morally permissible is that it is not one that a perfectly virtuous person would characteristically refrain from doing in the circumstances. I borrow the terms 'knowledge-supposing sense' and 'evidence-relative sense' from Parfit forthcoming, although the definitions are mine.

²⁶ Since there are different types of theories regarding what we ought to do (e.g., theories of morality, rationality, prudence, etc.), we'll need to add some modifier (e.g., morally, rationally, prudentially, etc.) in front of the word 'permissible' to restrict it to the relevant domain. And we'll need to insert the equivalent modifier (e.g., moral, rational, prudential, etc.) in the phrase 'better object-given reasons' so as to restrict the class of reasons to the same domain. Also, I should note that PMV is admittedly a rather weak claim, for even a theory that permits all actions embodies PMV. (I thank Jamie Dreier for making this clear to me.) Nevertheless, a theory that permits all actions does have one thing going for it: at least, it doesn't implausibly hold that it is sometimes impermissible for an agent to act so as to bring about the outcome that she ought to prefer above all other available alternatives.

bring about a lower ranked outcome instead of a higher ranked outcome.²⁷ PMV is to be preferred to *a*, because, given Schroeder's arguments and *a*, we would have to deny that other consequentialist theories, such as ethical egoism and self/other utilitarianism, share with act-utilitarianism the same compelling idea. And PMV is to be preferred to *b*, because *b* would have us deny that satisficing act-utilitarianism shares the same compelling idea that lies at the heart of maximizing act-utilitarianism.

To better understand why we should find PMV so compelling, we need only note how puzzling a theory that denies PMV can be in that it allows for the possibility that an agent ought to prefer an outcome in which she acts wrongly to any alternative in which she acts permissibly. Why this is so and why this should seem troubling is something that I will explain presently. For the purposes of illustration, I'll be using rule-utilitarianism as a representative view that denies PMV. Rule-utilitarianism makes for a particularly apt target, since it's clear what its axiological commitments are. Keep in mind that rule-utilitarianism is a *nonconsequentialist* theory given my convention of using the term 'consequentialism' as shorthand for 'maximizing *act*-consequentialism'.

To illustrate why it might seem troubling to deny PMV and thereby allow for the possibility that an agent ought to prefer the outcome in which she acts wrongly to any alternative in which she acts permissibly, consider the following example. Suppose that Jane must choose to do either *a*₁ or *a*₂, and that Jane's doing *a*₁ will entail her violating the ideal code, but that Jane's doing *a*₂ will entail that she does not violate the ideal code—the ideal code being the set of rules whose internalization by the overwhelming majority has maximum expected value. And assume that there is greater utility in *o*₁ than in *o*₂.

Now imagine a rule-utilitarian named Richard. Being a rule-utilitarian, Richard must accept welfarism, the view that an agent always has better object-given reasons to prefer an outcome with more utility to an outcome with less utility, even when she must violate the ideal code in order to bring about the outcome with more utility.²⁸ Thus Richard must accept that

²⁷ Claim *b* is what Foot seems to have in mind when she identifies the compelling idea as the "rather simple thought that it can never be right to prefer a worse state of affairs to a better" (1985, 198).

²⁸ Of course, this is not how welfarism is usually understood, but it seems to me to be a plausible way of understanding the view. In any case, the point here is to illustrate what's puzzling about the denial of PMV. So if one objects to this characterization of welfarism, then just assume that Richard is a rule-utilitarian who further accepts the buck-passing account of value.

Jane has better object-given reasons to prefer o_1 to o_2 than to prefer o_2 to o_1 . And since what an agent ought to prefer is simply what she has best object-given reasons to prefer, it follows that Richard must hold that Jane ought to prefer o_1 to o_2 . That what an agent ought to prefer is simply what she has best object-reasons to prefer isn't obvious, so let me pause to defend it. To see that this is so, consider the case where an all-powerful and all-knowing evil demon has threatened to torture me at noon tomorrow unless, by then, I both believe that $2 + 2 = 1$ and intrinsically desire to be tortured.²⁹ Given the evil demon's threat, it does seem that I ought to want to have this belief and desire and that I ought to do what I can to bring it about that I have this belief and desire by noon tomorrow. Yet it seems false to say that I ought either to believe that $2 + 2 = 1$ or to intrinsically desire to be tortured. Despite my knowledge of the evil demon's threat, it would be irrational for me either to believe that $2 + 2 = 1$ or to intrinsically desire to be tortured. This is because I cannot directly respond to the state-given reasons that the demon's threat provides me with. I can directly respond only to the object-given reasons I have both for wanting to be in those states and for causing myself to be in them if I can (Parfit forthcoming). If this is unconvincing, then just assume for the purposes of illustration that Jane has no state-given reasons to prefer o_1 to o_2 . In that case, given the presence of object-given reasons to prefer o_1 to o_2 and the lack of any state-given reasons not to, Richard must hold that Jane ought to prefer o_1 to o_2 . Either way, Richard, the welfarist, is committed to the view that Jane ought to prefer o_1 to o_2 .

Being a rule-utilitarian commits Richard not only to welfarism but also to rule-consequentialism. And, as a rule-consequentialist, Richard is committed to the view that Jane ought to refrain from doing a_1 and do a_2 instead, as this is the only way for Jane to avoid violating the ideal code. Therefore, Richard is committed both to the claim that Jane ought to prefer o_1 to o_2 and to the claim that Jane ought to do a_2 as opposed to a_1 . This puts Jane in a rather awkward position, because this means that if she responds to the situation as she ought to, she will do a_2 and then regret having done so, for it is plausible to suppose that if Jane ought to prefer the outcome where she has done a_1 as opposed to a_2 , then she ought to regret having done a_2 as opposed to a_1 . There is something quite strange about a theory that instructs an agent to do something and then regret having done so, and yet this is precisely what a theory that denies PMV does.

²⁹ This sort of example comes from Parfit forthcoming.

Also troubling is the fact that Richard's two claims—that Jane ought to prefer o_1 to o_2 and that Jane ought to do a_2 as opposed to a_1 —turn out to yield a contradiction on the teleological view of practical reasons. According to that view, “since any rational action must aim at some result, reasons that bear on whether to perform an action must appeal to the desirability or undesirability of having that result occur, taking into account the intrinsic value of the action itself” (Scanlon 1998, 84). On this view, then, if Jane has better object-given reasons to prefer o_1 to o_2 than to prefer o_2 to o_1 , then Jane has better reason to do a_1 than to do a_2 . Since Richard is committed to welfarism and thus to the view that Jane has better object-given reasons to prefer o_1 to o_2 , it follows, on the teleological view, that Jane has better reason to do a_1 than to do a_2 . And since a_1 and a_2 are the only alternatives, we must conclude that Jane ought to do a_1 , that which she has best reason to do. But this conflicts with Richard's claim that Jane ought to do a_2 , not a_1 . We see, then, that given the teleological view, Richard's two claims imply a contradiction.

Of course, one can reject the teleological view and argue both that Jane has non-teleological reasons to refrain from doing a_1 (perhaps, the fact that doing a_1 will violate the ideal code provides such a reason) and that these reasons outweigh, trump, undermine, or exclude the teleological reasons that Jane has to do a_1 . In that case, Richard could deny that Jane's having better object-given reasons to prefer o_1 to o_2 implies that Jane has better reason to do a_1 than to do a_2 . So one can deny PMV without contradiction but only if one denies the teleological view as well. Of course, it will come as no surprise that consequentialists find the teleological view attractive, and, if they were the only ones, we might question whether PMV is what nonconsequentialists find so compelling about consequentialism. Yet, as even Thomas Scanlon notes, many who reject consequentialism as a theory of morality find the teleological view of practical reasons attractive (1998, 81), and these nonconsequentialists will have to accept PMV to avoid contradiction.³⁰ Furthermore, the teleological view is still very much a major contender despite Scanlon's arguments against it—see, for instance, Arneson 2002 and Hurka forthcoming for how the teleologist can respond to Scanlon's challenges.

I've have provided two reasons for thinking that PMV is compelling. First, the denial of PMV can seem implausible once we see that it forces us

³⁰ It is possible to accept PMV and remain nonconsequentialist. To do so, one must hold that evaluative judgments (judgments concerning how alternative outcomes rank) are entirely dependent upon deontic judgments (judgments concerning the deontic statuses of the acts that produce them).

to accept the possibility that an agent might be required to do something and then regret having done so. Second, the denial of PMV leads to a contradiction on the teleological view of practical reasons, a view that many philosophers, including many nonconsequentialists, find attractive. In any case, if we do accept the compellingness of PMV, we will have a good explanation for why the move from act-consequentialism to rule-consequentialism has seemed such an unattractive solution to the problem of reconciling consequentialism with our commonsense moral convictions: in order to accommodate our convictions, rule-consequentialism must give up the very thing that we found to be most compelling about consequentialism in the first place, that is, the idea that it is never wrong for an agent to act so as to bring about the outcome that she ought to prefer above all other available alternatives. We would also have a good explanation for why act-utilitarianism tends to haunt even those who refuse to accept it. If this is right, then, despite Schroeder's arguments to the contrary, the consequentializing project is well motivated, for all consequentialist theories, as broadly defined, accept PMV—the determination condition ensures that they do. Agent-relative consequentialism can, then, endorse PMV and at the same time appeal to the good-relative-to relation to accommodate agent-centered constraints. Thus agent-relative consequentialism promises to take what's best from act-utilitarianism (namely, PMV) while leaving behind its counter-intuitive implications, at least, those stemming from its inability to accommodate agent-centered constraints. Of course, act-utilitarianism's inability to accommodate agent-centered constraints is only one source of its counter-intuitive implications. Another source lies with its inability to accommodate agent-centered options and supererogatory acts.³¹ Thus the consequentializer must find a way to accommodate these as well, and so, in the next section, I will show how they might do so.

4. The Deontic Equivalence Thesis

A number of philosophers (e.g., Broome 1991; Brown 2005; and Dreier 1993) have argued that, for any remotely plausible nonconsequentialist theory M , there is a consequentialist counterpart theory M^* such that the

³¹ Agent-centered options are options to do something other than what would best promote the impersonal good. They include *agent-favoring options*—options to favor one's own interests above those of others—as well as *agent-sacrificing options*—options to forgo some benefit for oneself in order to provide others with some lesser net benefit. The distinction between these two kinds of agent-centered options comes from Slote 1985.

set of actions that are permissible according to M^* is coextensive with the set of actions that are permissible according to M . These arguments, however, have fallen short in at least two respects. First, because these philosophers have focused on demonstrating only that, for any M , there is an M^* that mimics perfectly M 's permissibility verdicts, they have failed to demonstrate that, for any M , there is an M^* that mimics perfectly *all* of M 's deontic verdicts, including not only such verdicts as 'permissible' and 'impermissible', but also such verdicts as 'supererogatory'.³² Second, because these philosophers have been exclusively preoccupied with showing how a consequentialist theory can accommodate agent-centered constraints, they have failed to address what is perhaps the greatest challenge for the consequentializer: accommodating agent-centered options. Related to both these points is whether satisficing consequentialism is tenable, for perhaps the explanation for why these philosophers have neglected to explain how consequentializing might work with regard to both agent-centered options and supererogatory acts lies with the fact that satisficing consequentialism can easily accommodate both.³³ Unfortunately, though, there is good reason to question the tenability of satisficing consequentialism.³⁴ And if satisficing consequentialism is indeed untenable, then all hope for the success of the consequentializing project will lie with consequentialism (i.e., maximizing act-consequentialism). So, to make up for the shortcomings of my predecessors, I will argue for a considerably more ambitious thesis, the deontic equivalence thesis:

³² For instance, Brown argues for what he calls "Dreier's Conjecture: For any plausible moral theory M , there is some conceivable theory of the good G such that the set of actions that are right [i.e., permissible] according to M is coextensive with the set of actions whose outcomes are best according to G " (2004, 6). But, clearly, this does not establish that the resulting counterpart consequentialist theory will mimic *all* of M 's deontic verdicts, including its verdicts about which acts are supererogatory. Nevertheless, Brown does propose a redefinition of 'supererogatory acts', where consequentialism could, then, accommodate what he labels "supererogatory acts" (2004, 33-34). This, however, doesn't show that consequentialism can accommodate supererogatory acts, as ordinarily conceived, but only that consequentialism can accommodate what Brown has mislabeled "supererogatory acts."

³³ However, this isn't the correct explanation in Brown's case. Indeed, he's inclined to think that "we should reject the notions of satisficing and supererogation as incoherent" (2004, 33).

³⁴ My deepest worry lies with the fact that satisficing consequentialism implausibly permits agents to bring about a sub-optimal outcome (provided that sub-optimal outcome is good enough) merely for the sake of preventing a better outcome from obtaining. See Bradley forthcoming.

DET For any remotely plausible nonconsequentialist theory, M , there is a consequentialist counterpart theory, M^* , that yields, in every possible world, the exact same set of deontic verdicts that M does, *including not only such verdicts as ‘permissible’ and ‘impermissible’, but also such verdicts as ‘supererogatory’.*

As you will recall, the recipe for consequentializing a nonconsequentialist theory is simple. Take whatever considerations that the nonconsequentialist theory holds to be relevant to determining the deontic status of an action and insist that they are considerations relevant to ranking outcomes. So if M holds that it is wrong for an agent to dirty her hands and violate someone’s rights even for the sake of minimizing rights violations overall, then M^* need only hold that an outcome where P violates someone’s rights ranks lower on the $\succ_{v/t}$ ordering than any outcome where P doesn’t do so, even if P ’s violating someone’s rights would prevent others from committing more numerous rights violations. Furthermore, if M holds that it’s wrong to violate someone’s rights now even in order to prevent oneself from committing more numerous violations of that right in the future, then M^* need only hold that the outcome where P violates someone’s rights now ranks lower on the $\succ_{v/t}$ ordering than any outcome where P doesn’t do so, even if P ’s violating someone’s rights now would prevent P from committing more numerous violations of that right in the future. In like fashion, M^* could accommodate any other agent-centered constraint that M incorporates.

The two biggest issues regarding DET, though, are whether consequentialism can accommodate agent-centered options and whether consequentialism can accommodate supererogatory acts. Before addressing these issues, however, it is important to understand what both supererogatory acts and agent-centered options are and why it might appear impossible for consequentialism to accommodate them. I’ll start with supererogation.

A supererogatory act is, as the name suggests, an act that goes above and beyond what’s required. As Michael Byron has pointed out, “It’s difficult to see how a maximizing conception of morality can allow room for supererogation: If I’m required in every case to choose the best available option, how could I ever do *more* than morality requires?” (Byron 2004, 9). The answer depends on what doing *more* than is required consists in. If doing more than is required consists in merely doing more for others than one is required to do (whether or not there is any moral reason to do more for others than is required), then even traditional act-utilitarianism

can accommodate supererogatory acts. Traditional act-utilitarianism is the maximizing view that claims the following:

TAU An act is permissible just in case it produces at least as much aggregate utility as any other available alternative act would.

To illustrate, suppose that we were to define a supererogatory act as one that involves a greater self-sacrifice for the sake of others than is required.³⁵ In that case, TAU will accommodate supererogatory acts, because sometimes an agent will have a choice between two or more optimal acts that differ with respect to whether the optimal quantity of utility is distributed in greater proportion to others or in greater proportion to the agent. Such equally optimal alternatives will be morally optional on TAU, and those which involve *more* of the optimal quantity of utility being distributed to others as opposed to the agent will be supererogatory—at least, on the above definition.

But we might wonder whether this definition gets it right, specifically whether this is the relevant sense of ‘doing more than is required’. In specifying that ‘doing more than is required’ entails ‘doing more for others than is required’, this definition rules out the possibility of supererogation with respect to self-regarding duties.³⁶ Yet it certainly seems possible to go above and beyond what such duties require. For instance, we might think that there is a duty to develop one’s talents and that this is an imperfect duty, one that doesn’t require us to take every opportunity to develop our talents nor to develop our talents to the greatest extent possible, but that requires us to develop our talents only to a certain extent, taking advantage of a sufficient number of opportunities to develop our talents so as to meet the required threshold. But if this is right, then it is surely possible to go above and beyond one’s duty to develop one’s talents, and yet doing so may be of no benefit to others. Here, then, is a plausible candidate for supererogation that the above definition simply rules out of court.³⁷

This definition faces another, potentially more serious, problem. Note that, in the case of supererogation, it clearly matters what one is doing

³⁵ Sterling Harwood (2003, 182) defines supererogation in this way and, consequently, concludes that TAU can accommodate supererogatory acts.

³⁶ Jason Kawall makes this point here: <http://peasoup.typepad.com/peasoup/2004/09/supererogation_.html#comments>.

³⁷ If we think that supererogatory acts have to involve self-sacrifice, then we need only imagine a case where developing one’s talents beyond the extent to which one is required to will involve self-sacrifice.

more of. A supererogatory act couldn't be one that merely involves more perspiration than is required, as if taking a more circuitous route to achieving some moral good could be supererogatory merely because it would involve perspiring more than one is required. It seems, then, that a supererogatory act must involve doing more of whatever there is moral reason to do more of.³⁸ Since we have no moral reason to perspire more than we have to, performing an act that involves perspiring more than is required is not supererogatory. After all, it would be odd to allow that an act can be supererogatory even if the agent has no more moral reason to perform it than to perform some non-supererogatory alternative. Thus we should reject the above definition and accept instead the following: an act is supererogatory only if there is more moral reason to perform it than to perform some other permissible alternative.³⁹

If this is correct, then TAU cannot accommodate supererogatory acts, for, on TAU, the only time one has an option between, say, performing either a_1 or a_2 is when there is no more moral reason to perform a_1 than to perform a_2 and vice versa. On TAU, there is no moral reason to prefer that a given quantity of utility be distributed in greater proportion to others than to oneself; how utility is distributed is, on TAU, morally irrelevant. So, on TAU, agents can never do more than they are required to do in the relevant sense of 'doing more'. It remains unclear, then, how a consequentialist theory might incorporate supererogatory acts. Moreover, in regards to DET, it doesn't really matter whether what I say about supererogation is correct. If it isn't, we could still invent a name for an act that has the deontic status of being an act that the agent has more moral reason to perform than some permissible alternative and wonder whether consequentialism could mimic a theory that ascribed such a status to certain actions.

³⁸ Moral reasons are, of course, a proper subset of reasons for action; specifically, moral reasons are those reasons that can give rise to a moral ought, where 'ought' is understood broadly to express either obligation or advisability. Thus moral reasons are reasons that can give rise to an act's being either morally obligatory or morally supererogatory. I acknowledge that this sort of circular definition won't, by itself, allow us to differentiate moral reasons from non-moral reasons, but, at any rate, I take that to be the job of a substantive moral theory, not some definition. Different substantive moral theories will give different accounts of what sorts of facts constitute moral reasons.

³⁹ I leave open the question of whether there are any further necessary conditions, such as "an act is supererogatory only if it requires more self-sacrifice (or is otherwise more difficult or demanding to perform) than some other permissible alternative that one has less moral reason to do."

Note also that although TAU can accommodate morally optional acts (i.e., acts that one may both permissibly perform and permissibly refrain from performing), TAU cannot accommodate agent-centered options. Agent-centered options are options to do less than one can to promote the impersonal good. TAU can accommodate morally optional acts since there is sometimes more than one act that will be tied for first place in terms of utility production. But since it is only when two or more acts are tied for first place in terms of their utility production that TAU allows for an act's being morally optional, TAU cannot accommodate agent-centered options, options to do less than one can to promote utility. So not only do we need to show that maximizing consequentialism can incorporate supererogation, but we also need to show that it can incorporate agent-centered options.

Let me start with agent-centered options. To consequentialize a nonconsequentialist theory that incorporates agent-centered options, we will need to hold that whatever factor nonconsequentialist theories take to be relevant to determining whether doing one's best to promote the impersonal good is optional or required is a factor that's relevant to determining how outcomes are to be ranked. Arguably, what allows nonconsequentialist theories to incorporate agent-centered options is the assertion that how much cost there is to the agent in maximizing the impersonal good is relevant to determining whether refraining from doing so is permissible or not.⁴⁰ In that case, the consequentializer needs to find a way of ranking outcomes in terms of how much cost there is to the agent in her maximizing the impersonal good. To do so, the consequentializer need only define the $\succ_{v/t}$ ordering as follows:

SU Schefflerian Utilitarianism: $o_i \succ_{v/t} o_j \leftrightarrow$ both (1) $U_{-P}(o_i) > U_{-P}(o_j)$ and (2) $U_{+P}(o_i) \geq U_{+P}(o_j)$, where $U(o_i)$ denotes what the total aggregate utility (across time and individuals) would be were o_i to obtain, where $U_P(o_i)$ denotes what P 's lifetime utility would be were o_i to obtain, where $U_{-P}(o_i) = U(o_i) - U_P(o_i)$, and where $U_{+P}(o_i) = U_{-P}(o_i) + [10 \times U_P(o_i)]$.⁴¹

⁴⁰ Kagan calls this defense of options "the appeal to cost"—see his 1989.

⁴¹ I call this Schefflerian Utilitarianism, because like the theory that Scheffler argues for in his 1982 it incorporates agent-centered options (or what Scheffler calls 'agent-centered prerogatives'), but eschews agent-centered constraints (or what Scheffler calls 'agent-centered restrictions'). SU does, however, differ from Scheffler's Distributive Hybrid Theory in at least two important respects: (1) it doesn't give priority to benefiting the worse off and (2) it incorporates not only agent-favoring options, but also agent-sacrificing options.

To put things in plainer English, SU holds that P's act is permissible if and only if there is no available alternative act whose outcome contains both (1) more total utility for others (i.e., for those other than P) and (2) at least as much total adjusted utility, where we include everyone's utility but adjust the overall total by giving P's utility ten times the weight of anyone else's. Unlike TAU, SU accommodates many of the basic features of commonsense morality—specifically, supererogation, agent-centered options, and the self-other asymmetry.⁴² To illustrate, consider the case below, where an agent named Jill has the following four mutually exclusive and jointly exhaustive options:

a_i	$U(o_i)$	$U_P(o_i)$	$U_{-P}(o_i)$	$U_{+P}(o_i)$	Deontic Status of a_i
a_1	17	7	10	80	merely permissible
a_2	7	2	5	25	impermissible
a_3	16	1	15	25	supererogatory
a_4	19	-1	20	10	supererogatory

SU accommodates agent-centered options in that it permits agents to give their own utility anywhere from no weight at all up to ten times the weight of any other. Thus Jill can permissibly choose to perform a_1 as opposed to a_4 even though performing the latter would do more to promote the impersonal good.⁴³ Furthermore, by allowing agents to give their own utility no weight, SU even permits agents to forgo their own greater benefit for the sake of providing still others with some lesser net benefit. So, for instance, Jill can permissibly choose to perform a_3 as opposed to a_1 , thereby providing others with a net benefit of five utils at a cost of six utils to herself.⁴⁴ Thus SU accommodates what is called the 'self-other asymmetry'. On commonsense morality, there is an asymmetry between what an agent is permitted to do to herself and what she is permitted to do to others (hence, the name): specifically, whereas it is permissible for an agent to sacrifice her own greater benefit for the sake of providing others with some lesser net benefit, it is not permissible for an agent to sacrifice someone else's greater benefit (even with her permission) for the sake of providing others with some lesser net benefit.⁴⁵ Of course, the self-other asymmetry

⁴² Unlike Theodore Sider's "Self/Other Utilitarianism," SU permits agents not only to give their own utility no weight but, alternatively, to give it extra weight—see his 1993.

⁴³ $U(o_1)$ is 17, whereas $U(o_4)$ is 19.

⁴⁴ $U_{-P}(o_3)$ minus $U_{-P}(o_1)$ equals 5, whereas $U_P(o_3)$ minus $U_P(o_1)$ equals -6.

⁴⁵ See Slote 1985, chap. 1.

doesn't imply that it is always permissible to make self-sacrifices, only that it is permissible to make self-sacrifices that don't make things worse for others overall. Thus it would be wrong, as SU implies, for Jill to choose to perform a_2 as opposed to a_1 . Such a choice would not only be worse for Jill, but also worse for others as well.⁴⁶

Lastly, SU can accommodate supererogatory acts provided we assume that, on SU, agents have more moral reason to increase the utility of others than to increase their own utility.⁴⁷ In that case, Jill would have more moral reason to perform either a_3 or a_4 than she has to perform a_1 , a permissible alternative.⁴⁸ Moreover, SU allows for a range of supererogatory alternatives, where some supererogatory alternatives involve going further beyond what's required than others.⁴⁹ For instance, although both a_3 and a_4 are supererogatory, a_4 involves going further beyond what's morally required of Jill than a_3 does.

Of course, SU, as is, doesn't incorporate agent-centered constraints, but, as we've already seen, the process by which we might modify SU to arrive at a consequentialist theory that did incorporate agent-centered constraints is rather straightforward.⁵⁰ Thus, we can, I think, safely conclude that DET

⁴⁶ Whereas $U_P(o_1)$ is 7, $U_P(o_2)$ is only 2. And whereas $U_{-P}(o_1)$ is 10, $U_{-P}(o_2)$ is only 5.

⁴⁷ This assumption would explain why it is, on SU, always morally permissible to do what will be best for others overall, for, arguably, it is always morally permissible to do what one has most moral reason to do. For more on this point, see my 2005a. In that article, I argue that it is both always permissible to do what one has most moral reason to do and always permissible to do what one has at least as much reason, all things considered, to do. Thus, we might best understand the first conjunct in SU [$U_{-P}(o_i) > U_{-P}(o_j)$] to be specifying what it is that P has most moral reason to do and understand the second conjunct in SU [$U_{+P}(o_i) \geq U_{+P}(o_j)$] to be specifying what it is that P has at least as much reason, all things considered, to do.

⁴⁸ $U_{-P}(o_3)$ and $U_{-P}(o_4)$ are 15 and 20, respectively, whereas $U_{-P}(o_1)$ is only 10.

⁴⁹ On commonsense morality, there is often a range of supererogatory acts, where some supererogatory acts go even further beyond what's required than other supererogatory acts. For instance, it seems that not only is my choosing to spend my Saturdays helping the poor supererogatory, but so is my choosing to spend both my Saturdays and my Sundays helping the poor. Although both choices would seem to be supererogatory, the latter seems to go even further beyond what's required of me than the former does.

⁵⁰ One very crude way of modifying SU so as to accommodate both restrictions and special obligations goes as follows:

MSU Modified Schefflerian Utilitarianism: $O_i \succ_{P/t} O_j \leftrightarrow$ both (1) $U_{-P^*}(o_i) > U_{-P^*}(o_j)$ and (2) $U_{+P^*}(o_i) \geq U_{+P^*}(o_j)$, where $U_P(o_i)$ denotes what P's lifetime utility would be were o_i to obtain, where $U_{adj}(o_i)$ denotes what the total adjusted aggregate utility would be were o_i to obtain—adjusted both by multiplying any disutility due to harms that P caused by a factor of 1000 and by multiplying any utility due to P's

is true. Although we haven't considered every conceivable plausible nonconsequentialist theory and shown that there is a consequentialist counterpart theory that mimics all of its deontic verdicts, we have seen that a consequentialist theory can accommodate most, if not all, of the sorts of features a plausible nonconsequentialist theory might incorporate: supererogation, agent-centered options, agent-centered constraints, and the self-other asymmetry.⁵¹

5. What doesn't follow from the Deontic Equivalence Thesis

In this section, I respond both to those (e.g., Louise 2004) who have argued (*i*) that, from DET, it follows that we are all consequentialists and to those (e.g., Brown 2004) who have argued (*ii*) that, from DET, it follows that the consequentialism/nonconsequentialism distinction is empty. I argue that *i* can't be right since analogues of DET are true of most moral theories, including contractualism, Kantianism, virtue theory, etc. Thus if DET establishes that we're all consequentialists, then these analogues establish that we are all Kantians, contractualists, and virtue theorists as well, and all at the same time, which is just absurd. Against *ii*, I argue that moral theories are in the business of doing a lot more than just identifying which acts are right and which acts are wrong. So even if two moral theories are extensionally equivalent in their deontic verdicts, there can still be something substantive at issue between them. Lastly, I consider more recent attempts to argue that from DET and an additional premise (viz., what has been called "Foot's Thesis"), it follows that the consequentialism/nonconsequentialism distinction is empty. I argue that this can't prove that consequentialism is empty, for adopting Foot's Thesis amounts to abandoning consequentialism.

Against *i*, I will argue that there are analogues of DET for most, if not all, moral theories. Take Kantianism, for instance. The Kantian analogue of DET is: for any remotely plausible non-Kantian theory, M_{-K} , there is a Kantian counterpart theory, M_K , that yields, in every possible world, the exact same set of deontic verdicts that M_{-K} does. To illustrate, suppose that

fulfillment of a special obligation by 500, where $U_{-P^s}(o_i) = U_{adj}(o_i) - U_P(o_i)$, and where $U_{+P^s}(o_i) = U_{-P^s}(o_i) + [10 \times U_P(o_i)]$.

⁵¹ I haven't shown that the consequentializer can accommodate genuine moral dilemmas, cases where no matter what an agent does in her given choice situation she will be doing something morally impermissible. Of course, whether a theory that incorporates moral dilemmas would count as a plausible theory is contentious. But see Brown 2004 for an explanation of how the consequentializer might accommodate moral dilemmas.

a Kantian takes the principle “an act is morally permissible if and only if it involves treating humanity as an end-in-itself and never simply as a means” to be the one and only fundamental moral principle. The recipe for Kantianizing is, then, as follows. Take whatever considerations M_K regards as relevant to determining the deontic status of an action and insist that those considerations are relevant to determining whether humanity has been treated as an end-in-itself. So suppose that M_K is traditional act-utilitarianism (TAU). The theorist who wants to Kantianize TAU need only insist that we treat humanity as an end-in-itself if and only if we give equal consideration to everyone’s interests in maximizing aggregate utility. Similarly, we could follow the same procedure for any other M_K .

Like the traditional act-utilitarian, the utilitarian Kantian (the one who has Kantianized TAU) will hold that an act is morally permissible if and only if it maximizes aggregate utility; the difference is that the former, but not the latter, takes this to be a *fundamental* moral principle—a moral principle that does not derive from any more general moral principle. So whereas the traditional act-utilitarian believes that what ultimately makes an act wrong is that it fails to maximize aggregate utility, the utilitarian Kantian believes that what ultimately makes an act wrong is that it fails to treat humanity as an end-in-itself.

To take another example, consider how we could contractualize TAU. Suppose that a contractualist holds, as a fundamental moral principle, that an act is right if and only if it accords with principles “that no one could reasonably reject as a basis for informed, unforced general agreement” (Scanlon 1998, 153). In order to contractualize TAU, we need only insist that the only principle “that no one could reasonably reject as a basis for informed, unforced general agreement” is the principle of utility. And, again, we can follow the same procedure for contractualizing any other M_C . And so it seems we should accept the contractualist analogue of DET: for any remotely plausible non-contractualist theory, M_C , there is a contractualist counterpart theory, M_C , that yields, in every possible world, the exact same set of deontic verdicts that M_C does.

From these two illustrations, it should be evident how we might succeed in constructing analogues of DET for almost any moral theory. But if we are to assume that from DET, and DET alone, it follows that we’re all consequentialists, then, by the same reasoning, it follows from these analogues that we are also all contractualists and all Kantians. But this is patently absurd, precisely because these are *rival* moral theories that accept different fundamental moral principles and, hence, different rationales for the deontic verdicts that they give.

Others (e.g., Brown 2004) are not so much concerned to argue that we are all consequentialists as to argue that consequentialism is empty. Of course, to conclude that consequentialism is empty from the mere fact that DET is true (which, admittedly, Brown doesn't do) is to presume that if two moral theories are necessarily coextensive in their deontic verdicts, then there is nothing substantive at issue between them. But this is false, because moral theories do much more than just yield moral verdicts. Importantly, they provide different competing rationales for the deontic verdicts that they yield. Thus, as we have just seen, an act-utilitarian, a Kantian, and a contractualist can all agree that the extension of permissible acts is just those that maximize utility, but even so they will provide different explanations for why this is so, for they necessarily accept different views about what the fundamental right-making and wrong-making features of acts are. And this means that the theories, although extensionally equivalent, will have different truth conditions. For instance, TAU is true only if utility is the only thing that is good for its own sake. And utilitarian contractualism is true only if the only principle that "no one could reasonably reject as a basis for informed, unforced general agreement" is the principle of utility. And utilitarian Kantianism is true only if giving everyone equal consideration in maximizing aggregate utility is what constitutes treating humanity as an end-in-itself. These are substantive issues over which the traditional act-utilitarian, the utilitarian contractualist, and the utilitarian Kantian may disagree.⁵²

In response to this, someone like Brown (2004) would argue that although *i* and *ii* may be false, we can conclude from DET and one additional premise that consequentialism is empty. That additional premise is what Brown calls:

FT Foot's Thesis: $o_i \succ_{p/t} o_j$ if and only if *P*'s doing a_i at *t* is deontically superior to *P*'s doing a_j at *t*, where one act is deontically superior to another if and only if the former is permissible and the latter is impermissible.⁵³

⁵² For more on this issue, see Portmore 2001.

⁵³ Obviously, this is not the way Foot would have formulated her thesis, nor is it the way that Brown formulates it. I've had to formulate it in terms of the $\succ_{p/t}$ ranking so that it dovetails with the rest of the paper. Nevertheless, the general idea remains the same: there is no way to make a judgment as to how outcomes are to be ranked "outside of morality," that is, independent of our judgments as to what's right and wrong. For the record, Brown formulates Foot's Thesis as follows: "A theory of the good is true if and only if the set of

A discussion of whether FT is plausible goes beyond the scope of this paper. Indeed, it goes beyond the scope of Brown's paper as well, for, as Brown admits, he provides no argument for FT (2004, 35)—and, for that matter, neither does Dreier (1993) or Louise (2004). Nevertheless, it's sufficient for my purposes to point out that FT amounts to the denial of what I've called the "determination condition" (see section 2 above) and hence to the denial of consequentialism. On FT, there are no independent criteria by which we can determine whether one outcome outranks another on the $\succ_{v/t}$ ordering. According to FT, the only criterion we can use to determine whether or not $o_i \succ_{v/t} o_j$ is the one where we ask whether or not P 's doing a_i at t is deontically superior to P 's doing a_j at t . But, in that case, the determination condition is false: acts can't have the deontic statuses that they do *in virtue* of how they rank if how they rank is just a matter of what their deontic statuses are. Thus FT amounts to the denial of consequentialism, and so it can't follow from the conjunction of FT and DET, that consequentialism is empty or that we're all consequentialists. Of course, 'consequentialism' is a term of art and Brown can define consequentialism, as he in fact does, such that consequentialism isn't tied to the determination condition and "doesn't say that actions are right *because* their outcomes are best" (2004, 4).⁵⁴ According to Brown, anyone who accepts the following bi-conditional is a consequentialist (2004, 4):

BC An act is permissible if and only if no other available outcome ranks higher, evaluatively speaking, than its outcome.

On this definition, Brown is, of course, correct in concluding that, from the conjunction of FT and DET, it follows that what Brown calls "consequentialism" (viz., BC) is empty. But note that on this definition Foot

actions whose outcomes are best according to that theory is coextensive with the set of actions that are right" (2004, 7).

Note that if the thought was merely that the only way to *know* whether a particular theory of the good is true or not is to plug it into BC and see whether the resulting deontic verdicts comport with our pre-theoretical intuitions, then this would establish that BC is uninformative, but not that it is empty—see section 6 below.

⁵⁴ Brown says, "consequentialism, as here defined, claims only that certain properties are coninstantiated.... Consequentialism so defined makes no claim about causation, determination, or justification..." (2004, 4). Yet later, on p. 26, he forgets himself and says, "A core idea of consequentialism is that rightness and wrongness are *determined* [his emphasis] by goodness." In correspondence, Brown has subsequently said that he would retract the latter statement.

is a “consequentialist.” Given her acceptance of FT, she would have no trouble accepting BC. This is, then, a cheap victory. If we are to define ‘consequentialism’ in such a way as to include such self-described nonconsequentialists as Foot, then the claim that we are all “consequentialists” (or that “consequentialism” is empty) loses much of its punch. That is, Brown has not established the surprising conclusion that we are all theorists of the type that Portmore and Frey profess to be and that Vallentyne and Kamm profess not to be, for these four all take the determination condition to be essential to consequentialism.⁵⁵ Thus Brown has not shown that we are all consequentialists or that consequentialism is empty as I and others have defined ‘consequentialism’.

6. Appealing to our moral convictions in specifying the relevant ordering

Since one of the consequentializer’s main aims is to avoid the counter-intuitive implications that plague TAU, the consequentializer will need to appeal to our considered moral convictions when determining how outcomes are to be ordered. But we might wonder whether this is legitimate. That is, we might wonder: Can a consequentialist defend one ordering over another on the basis that the former but not the latter yields intuitive moral verdicts when combined with BC? In this section, I argue that the consequentialist can legitimately do so and that in doing so consequentialism is rendered neither circular nor uninformative.⁵⁶

So suppose that someone accepts BC, that an act is permissible if and only if no other available outcome ranks higher, evaluatively speaking, than its outcome.⁵⁷ Having accepted BC, there are three possible procedures one might employ in figuring out the proper ranking of outcomes (i.e., figuring out how an agent ought to rank the various alternative outcomes that she’s capable of producing):

The Footian Procedure: Come to some fixed set of pre-theoretical judgments about the rightness of actions that is entirely independent of any judgments one has about the proper ranking of outcomes, and then

⁵⁵ Examples of those who include such a determination condition on their definition of ‘consequentialism’ abound—see, for instance, Bradley forthcoming, Frey 2000 (p. 165), Griffin 1995 (p. 154), Kamm 2000 (p. 205), and Vallentyne 2006.

⁵⁶ In conversation, Harry Silverstein raised this worry: the worry that, in doing so, consequentialism would be rendered circular and/or uninformative.

⁵⁷ From here on out, I’ll drop the qualifier ‘evaluatively speaking’, but the reader should take it to be implicit in the remaining.

hold that one outcome outranks another if and only if its corresponding act is deontically superior to that of the other.

The Foundationalist Procedure: Come to some fixed set of pre-theoretical judgments about the proper ranking of outcomes that is completely independent of any judgments one has about the rightness of actions, and then rank outcomes accordingly.

The Coherentist Procedure: Keeping BC constant, revise one's pre-theoretical judgments both about the proper ranking of outcomes and about the rightness of actions in light of each other until reflective equilibrium is reached, and then rank outcomes accordingly.

The Footian Procedure would not yield any normative advice beyond what our pre-theoretical deontic judgments already tell us. Thus the consequentialist cannot adopt the Footian Procedure if she wants the resulting substantive theory (the one obtained by applying the results of the Footian Procedure to BC) to yield informative results. Failing to provide any normative advice beyond what our pre-theoretical deontic judgments already tell us would be a significant failure, as one of the things that we hope to gain from moral theorizing is a deeper understanding of morality, so that we might better deal with moral questions about which we have no confident pre-theoretical judgment.⁵⁸

But just because the consequentializer should eschew the Footian Procedure doesn't mean that she must eschew any appeal to our pre-theoretical deontic judgments and adopt the Foundationalist Procedure. Fortunately for the consequentializer, the Coherentist Procedure is a viable alternative that, like the Foundationalist Procedure, yields informative results. Moreover, the Coherentist Procedure has the advantage of being likely to yield a substantive consequentialist theory (one that combines consequentialism with a substantive account of how outcomes are to be ordered) that avoids many of the counter-intuitive implications that plague TAU.

Unlike both the adherents of the Foundationalist and Footian Procedures, the adherent of the Coherentist Procedure doesn't consider either her pre-theoretical judgments about the proper ranking of outcomes or her pre-theoretical judgments about the rightness of actions to be fixed starting points. Whereas the adherent of the Foundational Procedure will

⁵⁸ See Hooker 2000, chap. 1.

be unwilling to revise her judgments about the proper ranking of outcomes in light of any potentially counter-intuitive implications it yields when combined with BC, the adherent of the Coherentist Procedure is willing to do so. And whereas the adherent of the Footian Procedure will be unwilling to revise her pre-theoretical judgments about the rightness of actions in light of any potential conflicts that may arise when she combines her pre-theoretical judgments about the proper ranking of outcomes with BC, the adherent of the Coherentist Procedure is willing to do so. Sometimes the adherent of the Coherentist Procedure will end up revising her deontic judgments in light of her more firmly held evaluative judgments. Other times she will revise her evaluative judgments in light of her more firmly held deontic judgments. The Coherentist Procedure will, therefore, be informative, yielding both new judgments about the deontic statuses of actions and new judgments about the how outcomes should be ranked. Given the viability of the Coherentist Procedure, we can conclude, then, that the consequentializer can appeal to our considered moral convictions when determining how outcomes are to be ordered without rendering her view either circular or uninformative.

7. Conclusion

The consequentializing project is to come up with a theory that takes what is best from both traditional act-utilitarianism and traditional nonconsequentialism, while leaving behind the disadvantages of each. From the former, the consequentializer hopes to take the compelling idea that it is always permissible to act so as to bring about the highest ranked available outcome (I've called this compelling idea "the permissibility-of-maximizing view" or "PMV.") From the latter, the consequentializer hopes to take its comportment with our considered moral convictions. The result is a very promising moral theory, a theory that (1) comports with our considered moral convictions, (2) unifies these convictions under a single moral principle, and (3) embodies the compelling idea that it is always permissible to bring about the outcome that one has better object-given reasons to prefer above all other available alternatives.⁵⁹ I've defended this

⁵⁹ In spite of all this, the nonconsequentialist can still argue that the consequentializer gets the right deontic verdicts but for the wrong reasons. That is, the nonconsequentialist can still claim that consequentialism is false, not because it ignores certain morally relevant factors, but because it provides the wrong rationale for why these factors are the morally relevant ones. But if PMV is indeed compelling and if, contrary to Foot, we can get some handle on how outcomes rank independent of the deontic statuses of the acts that produce

consequentializing project against a number of objections. Against Schroeder's criticisms, I've argued that the consequentializing project can succeed in simultaneously achieving its two aims. Against those who might argue that satisficing consequentialism is implausible and that there is no room for agent-centered options and supererogatory acts on a maximizing conception of morality, I've argued for DET: for any remotely plausible nonconsequentialist theory, *M*, there is a consequentialist (i.e., a maximizing act-consequentialist) counterpart theory, *M**, that yields, in every possible world, the exact same set of deontic verdicts that *M* does. And I've argued that DET doesn't render consequentialism empty. Lastly, against those that would argue the consequentializer's appeal to our considered moral convictions in determining how to rank outcomes threatens to render consequentialism circular or uninformative, I've argued that neither is the case.

Of course, there are still objections for the consequentializer to address. One objection, which comes from Foot, is that there is no way to rank outcomes outside of the scope of the sorts of comparisons of outcomes that the virtue of benevolence allows (Foot 1985).⁶⁰ Another more recent objection involves challenging the preeminence of the teleological view of practical reasons (Scanlon 1998).⁶¹ Both objections threaten to break the spell that consequentialism has had over us, the former by arguing that nonconsequentialists can accept PMV in so far as evaluative comparisons of outcomes make sense at all, the latter by arguing that the teleological view of practical reasons, on which the compellingness of PMV largely hinges, is flawed. Both would undercut the motivation for the consequentializing project. But these are challenges, not decisive objections to the consequentializing project, and, more importantly, they are topics for another paper.⁶²

them, then we have a compelling reason to think that the consequentialist's rationale is the more plausible one.

⁶⁰ Foot says, "...sometimes justice will forbid a certain action...and then it will not be possible to ask whether 'the state of affairs' containing that action and its results will be better or worse than one in which the action is done" (1985, 206). Nevertheless, others, like myself, remain unconvinced—see, for instance, Scheffler 1985, 412-413.

⁶¹ For some replies to Scanlon's objections to the teleological view of practical reasons, see, for instance, Arneson 2002 and Hurka forthcoming.

⁶² I've benefited immensely from a discussion of these issues with various commentators on *PEA Soup* <<http://peasoup.typepad.com/peasoup/>>, especially Ben Bradley, Campbell Brown, Jamie Dreier, Josh Glasgow, Robert Johnson, Troy Jollimore, Jason Kawall, Mark van Roojen, Jussi Suikkanen, and Kyle Swan. And I want to especially thank Peter de Marneffe, Jamie Dreier, Josh Glasgow, Jennie Louise, Mark Schroeder, Jussi Suikkanen,

*Department of Philosophy
Arizona State University
P.O. Box 874102
Tempe, AZ 85287-4102*

Douglas.Portmore@asu.edu

References

- Arenson, R. (2002). "The End of Welfare as We Know It? Scanlon versus Welfarist Consequentialism." *Social Theory and Practice* 28: 315-336.
- Bradley, B. (forthcoming). "Against Satisficing Consequentialism." *Utilitas*.
- Brentano, F. (1989). *The Origin of Our Knowledge of Right and Wrong*. Translated by R. M. Chisholm and E. H. Schneewind. London: Routledge and Kegan Paul, 1969.
- Broome, J. (1991). *Weighing Goods*. Oxford: Basil Blackwell.
- Brown, C. (2004). "Consequentialise This." Working manuscript of 9/10/04. <<http://personal.bgsu.edu/~browncf/papers/ConsequentialiseThis.pdf>>.
- Byron, M. (2004). "Introduction." In his (ed.) *Satisficing and Maximizing* (pp. 1-13). Cambridge: Cambridge University Press.
- Dreier, J. (1993). "Structures of Normative Theories." *The Monist* 76: 22-40.
- Ewing, A. C. (1947). *The Definition of Good*. London: Macmillan.
- Foot, P. (1985). "Utilitarianism and the Virtues." *Mind* 94: 196-209.
- Frey, R. G. (2000). "Act-Utilitarianism." In H. LaFollette (ed.), *The Blackwell Guide to Ethical Theory* (pp. 165-182). Oxford: Blackwell.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford: Clarendon Press.
- Griffin, J. (1995). "Consequentialism." In T. Honderich (ed.), *The Oxford Companion to Philosophy* (pp. 154-156). Oxford: Oxford University Press.
- Harwood, S. (1998). "Eleven Objections to Utilitarianism." In L. Pojman (ed.), *Moral Philosophy: A Reader* (pp. 179-192). Indianapolis: Hackett Publishing.
- Hooker, B. (2000). *Ideal Code, Real World*. Oxford: Clarendon Press.
- Hurka, T. (forthcoming). "Value and Friendship: A More Subtle View." *Utilitas*.
- . (1990). "Two Kinds of Satisficing." *Philosophical Studies* 59: 107-111.

Mark van Roojen, an anonymous referee from this journal, and an audience at Arizona State University for helpful comments and suggestions on an earlier draft of this paper.

- Kagan, S. (1989). *The Limits of Morality*. Oxford: Clarendon Press.
- . (1992). "The Structures of Normative Ethics." *Philosophical Perspectives* 6: 223-242.
- Kamm, F. M. (2000). "Nonconsequentialism." In H. LaFollette (ed.), *The Blackwell Guide to Ethical Theory* (pp. 165-182). Oxford: Blackwell.
- Louise, J. (2004). "Relativity of Value and the Consequentialist Umbrella." *Philosophical Quarterly* 54: 518-536.
- McDowell, J. (1985). "Values and Secondary Qualities." In T. Honderich (ed.), *Morality and Objectivity* (pp. 110-129). London: Routledge and Kegan Paul.
- Nozick, R. (1968). "Moral Complications and Moral Structures." *Natural Law Forum* 13: 1-50.
- . (1974). *Anarchy, State, and Utopia*. New York: Basic Books.
- Parfit, D. (forthcoming). *Climbing the Mountain*. Oxford: Oxford University Press.
- . (2001). "Rationality and Reasons." In D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmussen, eds., *Exploring Practical Philosophy*. Aldershot: Ashgate.
- . (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Pettit, P. (1997). "The Consequentialist Perspective." In M. Baron, P. Pettit, and M. Slote, *Three Methods of Ethics*. Oxford: Blackwell.
- Portmore, D. W. (2005a). "Dual-Ranking Act-Consequentialism." Working manuscript of 9/5/05. <http://www.public.asu.edu/~dportmor/Dual-Ranking_Act-Consequentialism.pdf>.
- . (2005b). "Combining Teleological Ethics with Evaluator Relativism: A Promising Result." *The Pacific Philosophical Quarterly* 86: 95-113.
- . (2003). "Position-Relative Consequentialism, Agent-Centered Options, and Supererogation." *Ethics* 113: 303-32.
- . (2001). "Can an Act-Consequentialist Theory be Agent-Relative?" *American Philosophical Quarterly* 38: 363-77.
- Post, J. F. (1999). "Is Supervenience Asymmetric?" In L. C. Pereira and M. Wrigley, eds., *Festschrift in Honor of Oswaldo Chateaubriand* (pp. 305-344). Manuscript.
- Rabinowicz, W. and Rønnow-Rasmussen, T. (2004). "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114: 391-423.
- Raz, J. (1975). *Practical Reasons and Norms*, (Princeton, NJ: Princeton University Press.
- Ridge, M. (Forthcoming). "Review of Pleasure and the Good Life." *Utilitas*.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, Mass.: The Belknap Press.

- Scheffler, S. (1982). *The Rejection of Consequentialism*. Oxford: Clarendon Press. Revised edition 1994.
- . (1985). "Agent-Centered Restrictions, Rationality, and the Virtues." *Mind* 94: 409-419. Reprinted in his edited volume *Consequentialism and Its Critics*. Oxford: Oxford University Press.
- Schroeder, M. (forthcoming). "No So Promising After All: Evaluator-Relative Teleology and Common-Sense Morality." *Pacific Philosophical Quarterly*.
- . (2005a). "Maximizing Agent-Relative Value." Working manuscript of 6/7/05. <<http://www.wam.umd.edu/%7Emschroed/Schroeder%20Maximizing%20Relative%20Value.pdf>>.
- . (2005b). "Teleology, Agent-Relative Value, and 'Good'." Working manuscript 5/29/05. <<http://www.wam.umd.edu/%7Emschroed/Schroeder%20Agent-Relative%20Teleology.pdf>>.
- Shaw, W. (2006). "The Consequentialist Perspective." In J. Dreier, ed., *Contemporary Debates in Moral Theory*. Oxford: Blackwell.
- Sider, T. (1993). "Asymmetry and Self-Sacrifice." *Philosophical Studies* 70: 117-132.
- Sinnott-Armstrong, W. (2003). "Consequentialism." *The Stanford Encyclopedia of Philosophy* (Summer 2003 Edition), Edward N. Zalta, ed., URL = <<http://plato.stanford.edu/archives/sum2003/entries/consequentialism/>>.
- Smith, M. (2003). "Neutral and Relative Value after Moore." *Ethics* 113: 576-598.
- Slote, M. (1985). *Common-Sense Morality and Consequentialism*. London: Routledge and Kegan Paul.
- . (1984). "Satisficing Consequentialism." *Proceedings of the Aristotelian Society* 58: 139-63.
- Sosa, D. (1993). "Consequences of Consequentialism." *Mind* 102: 101-22.
- Stratton-Lake, P. (2005). "How to Deal with Evil Demons: Comment on Rabinowicz and Rønnow-Rasmussen." *Ethics* 115: 788-798.
- Vallentyne, P. (forthcoming) "Against Maximizing Consequentialism." In J. Dreier, ed., *Contemporary Debates in Moral Theory*. Oxford: Blackwell.
- . (1988a). "Teleology, Consequentialism, and the Past." *The Journal of Value Inquiry* 22: 89-101.
- . (1988b). "Gimmicky Representation of Moral Theories." *Metaphilosophy* 19: 253-63.
- . (1987a). "The Teleological/Deontological Distinction." *Journal of Value Inquiry* 21: 21-32.
- . (1987b). "Utilitarianism and the Outcomes of Actions." *The Pacific*

Philosophical Quarterly 68: 57-70.

Wiggins, D. "A Sensible Subjectivism." In D Wiggins, *Needs, Value, Truth: Essays in the philosophy of value*. Oxford: Blackwell.

Williams, B. (1973). "A Critique of Utilitarianism." In J. J. C. Smart and B. Williams, *Utilitarianism: For and Against*. Cambridge: Cambridge University Press.