

Against Satisficing Consequentialism

BEN BRADLEY

Syracuse University

The move to satisficing has been thought to help consequentialists avoid the problem of demandingness. But this is a mistake. In this article I formulate several versions of satisficing consequentialism. I show that every version is unacceptable, because every version permits agents to bring about a submaximal outcome in order to prevent a better outcome from obtaining. Some satisficers try to avoid this problem by incorporating a notion of personal sacrifice into the view. I show that these attempts are unsuccessful. I conclude that, if satisficing consequentialism is to remain a position worth considering, satisficers must show (i) that the move to satisficing is necessary to solve some problem, whether it be the demandingness problem or some other problem, and (ii) that there is a version of the view that does not permit the gratuitous prevention of goodness.

Michael Slote famously suggested that consequentialists should make use of the notion of satisficing in order to bring consequentialism more in line with 'common sense'. Slote presented a number of examples designed to illustrate the permissibility of satisficing. It is permissible, he said, for the fairy-tale hero to ask the gods just to make him and his family comfortable, even though he could ask for much more, and for the motel owner to offer the stranded motorists the first available satisfactory room, rather than the best room in the motel.¹

While Slote has evidently abandoned satisficing consequentialism, other philosophers have taken up the cause, because of what they perceive to be the demandingness of maximizing consequentialism.² There are many ways to develop a satisficing consequentialist view. In what follows I formulate several versions of satisficing consequentialism. I show that every version is unacceptable, because

¹ M. Slote, *Common-Sense Morality and Consequentialism* (London, 1985), pp. 45–7.

² Peter Vallentyne claims that consequentialists ought to be satisficers (Vallentyne, 'Against Maximizing Act-Consequentialism', *Moral Theories*, ed. J. Dreier (forthcoming); Robert Elliot and Dale Jamieson seem to endorse something called 'improving consequentialism', which seems to be some sort of satisficing view (Elliot, *Faking Nature: The Ethics of Environmental Restoration* (New York, 1997), pp. 46–8; Jamieson, 'When Utilitarians Should Be Virtue Theorists', *Utilitas* (forthcoming); Thomas Hurka gives satisficing a qualified endorsement, claiming that 'absolute level' satisficing is plausible assuming a subjective theory of the good (Hurka, 'Two Kinds of Satisficing', *Philosophical Studies* 59 (1990), pp. 107–11). All seem to be motivated by the demandingness problem. Also see John Turri, 'You Can't Get Away With Murder That Easily: A Response to Tim Mulgan', *International Journal of Philosophical Studies* (forthcoming), and several of the entries in Michael Byron (ed.), *Satisficing and Maximizing: Moral Theorists on Practical Reason* (Cambridge, 2003).

every version permits agents to bring about a submaximal outcome in order to prevent a better outcome from obtaining.

I. VARIETIES OF SATISFICING CONSEQUENTIALISM

Satisficing consequentialists, in the sense relevant here, hold that an act is morally right if and only if its consequences are ‘good enough’. But this is too vague to be helpful. We need to know what ‘good enough’ means. There are many things it might mean.

Elaborating on remarks by Slote, Thomas Hurka points out a distinction between ‘absolute level’ and ‘comparative’ satisficing views.³ According to the absolute level view, there is a level of utility such that an act is right iff its utility reaches that level, no matter what the alternatives are. Of course, there might be situations in which that level cannot be reached by any available alternative. In those situations, Hurka suggests we should maximize utility. Here, then, is absolute level satisficing consequentialism:

ALSC There is a number, n , such that: An act is morally right iff either (i) it has a utility of at least n , or (ii) it maximizes utility.

According to the comparative view, one must bring about an outcome that is ‘reasonably close’ to the best outcome. There are at least two ways to cash out the notion of ‘reasonably close’. One way, suggested by Hurka, is to take a percentage.⁴ Here is a simple percentage view: take the best act available and multiply its value by some fraction to get the threshold that must be met to act rightly on that occasion. But this goes wrong in cases where even the best act has negative utility. A better idea is to look at the best alternative and the worst alternative, and take some percentage of the difference. In other words:

CSC1 There is a fraction, n ($0 < n < 1$), such that: An act is morally right iff its utility, plus [(the utility of a maximizing alternative – the utility of a minimizing alternative) multiplied by n], is at least as great as the utility of a maximizing alternative.

The smaller the fraction, the more demanding the view.⁵

³ Slote, *Common-Sense Morality*, pp. 50–1; Hurka, ‘Two Kinds of Satisficing’.

⁴ This is what Hurka suggests in the following passage: ‘no matter how bad the initial situation, [comparative satisficing] requires agents only to make some reasonable percentage of the largest improvement in it they can’ (‘Two Kinds of Satisficing’, 108).

⁵ Comparative satisficers might instead require an agent to perform an act that is at least as good as some percentile of the available alternatives (Vallentyne, ‘Against Maximizing’, p. 13, and Vallentyne, ‘Consequentialism’, *Ethics in Practice*, 3rd edn., ed. H. LaFollette (forthcoming), p. 12). This will not help solve the problems I point out below, so I will not discuss the idea further.

Here is a simpler way to understand ‘reasonably close’: in every circumstance, regardless of what the worst outcome is, one’s act must always have a utility that is no greater than a certain distance from the best outcome. More precisely:

CSC2 There is a number, n ($n > 0$), such that: An act is morally right iff its utility plus n is greater than or equal to the utility of a utility-maximizing alternative.

Finally, we can see that there is room to combine the absolute level view with the comparative view. Call this view ‘double satisficing’ consequentialism:

DSC There is a number, n , such that: An act is morally right iff either (i) it has a utility of at least n , or (ii) its utility is less than n , but reasonably close to that of a utility-maximizing alternative.⁶

Since we have seen that ‘reasonably close’ has two interpretations, there are two versions of DSC, one corresponding to CSC1 and the other corresponding to CSC2.

There is another important distinction to be made. According to the versions of satisficing consequentialism discussed so far, the permissibility of an act is determined by the value of what it produces. But on some satisficing views, what matters is the value found in a ‘situation’ or ‘state of affairs’ at or around the time of an act. This distinction needs further elaboration. When introducing the absolute level version of satisficing, Hurka characterizes the view in this way:

When a situation is and will remain below the absolute threshold, an agent’s duty is the same as under maximizing: she must do everything to move it towards satisfactory goodness. Once the threshold is reached, however, her duty vanishes. If a state of affairs is already, by absolute standards, reasonably good, she has no duty whatever to improve it.⁷

If the state of affairs before a person acts is below the absolute threshold, and cannot be raised to that threshold, absolute-level satisficing on its own tells him to make it as good as possible, or to bring it as close as possible to absolute satisfactoriness.⁸

For Hurka, what matters is the amount of value found in a *situation* before an act is performed. Hurka does not tell us what he means by a ‘situation’. Let us for now think of the situation before an act as consisting of everything that happens in the world in a very short period of time before the act is performed. The value of a situation is determined in some way by the values of the things that happen in it.

⁶ Hurka suggests that this is really Slote’s view (‘Two Kinds of Satisficing’, p. 109).

⁷ ‘Two Kinds of Satisficing’, p. 108.

⁸ *Ibid.*, p. 108.

Hurka suggests, in discussing the absolute level view, that if the situation before an act has at least a certain value, an agent need not do anything. But that cannot be right. Suppose that an agent finds himself in a situation with a value above the threshold level set by the absolute level view, and suppose that were he not to do anything, the situation after that inaction would have a value below the threshold. If what is important is the situation being above some threshold, then it seems obligatory to prevent the situation from falling below that threshold if possible. This suggests that what really matters is not the value of the situation *before* the act, but its value *after* the act. Here is how we might state an absolute level version of the 'situation' view:

SALSC: There is a number, n , such that: An act is morally right iff either (i) the situation that would obtain after the act has value of at least n , or (ii) the act maximizes utility.

SALSC does not seem to be just the view Hurka has in mind. 'If the good is pleasure, absolute-level satisficing requires agents to ensure that *everyone* is above some threshold of pleasure.'⁹ What is relevant for Hurka is not whether the world is above some value threshold at the time of an act; rather, it is whether each individual is above some threshold. Here is how we might state the individualist version of the situation version of the absolute level view:

ISALSC: There is a number, n , such that: An act is morally right iff either (i) in the situation after the act, each person's welfare level is at least n , or (ii) the act maximizes utility.

We could get different versions of satisficing consequentialism by giving the existential quantifier narrow scope. Rather than saying that there is some utility such that for any situation, an act performed in that situation must meet or exceed that utility, we might instead have said that for each situation, there is some utility an act must reach in order to be right – allowing the threshold to change depending on the situation. But to avoid turning satisficing consequentialism into a form of particularism, some story would have to be told about what features of the situation are relevant to determining the threshold. No such story has been told, and I am not interested in trying to tell it. So I will not consider such versions of satisficing consequentialism here.

⁹ Hurka, 'Satisficing and Substantive Values', in *Satisficing and Maximizing*, p. 72; my emphasis.

II. EVALUATIONS OF THE VARIETIES OF SC¹⁰

Just to recap, we now have the following versions of Satisficing Consequentialism to consider.

ALSC: There is a number, n , such that: An act is morally right iff either (i) it has a utility of at least n , or (ii) it maximizes utility.

CSC1: There is a fraction, n ($0 < n < 1$), such that: An act is morally right iff its utility, plus [(the utility of the best alternative – the utility of the worst alternative) multiplied by n], is at least as great as the utility of the best alternative.

CSC2: There is a number, n ($n > 0$), such that: An act is morally right iff its utility plus n is greater than or equal to the utility of a utility-maximizing alternative.

DSC: There is a number, n , such that: An act is morally right iff either (i) it has a utility of at least n , or (ii) its utility is less than n , but reasonably close to that of a utility-maximizing alternative.

SALSC: There is a number, n , such that: An act is morally right iff either (i) the situation that would obtain after the act has value of at least n , or (ii) the act maximizes utility.

ISALSC: There is a number, n , such that: An act is morally right iff either (i) in the situation after the act, each person's welfare level is at least n , or (ii) the act maximizes utility.

I will evaluate the 'situational' views first. SALSC is probably the least plausible version of SC. Suppose one happens to be in a situation whose value is far above the threshold level. One could permissibly kill several people gratuitously, so long as doing so would not cause the value of the situation to drop below the threshold.¹¹ This consequence will remain even if we (somehow) reformulate the situation view as a comparative view rather than an absolute level view.

ISALSC does not seem to have that very problem, since, for any plausible n , killing someone would likely cause his welfare level to fall below n .¹² But ISALSC has its own problems. One problem for ISALSC is that if n is reasonably high, there will always be someone below

¹⁰ Some of the examples in this section are similar, in certain important respects, to ones given by Tim Mulgan and Erik Carlson. See Mulgan, 'Slote's Satisficing Consequentialism', *Ratio* 6 (1993); Mulgan, 'How Satisficers Get Away with Murder', *International Journal of Philosophical Studies* 9 (2001); Mulgan, *The Demands of Consequentialism* (New York, 2001); and Carlson, *Consequentialism Reconsidered* (Dordrecht, 1995).

¹¹ Mulgan gives a similar argument, though his argument is directed against 'contribution' satisficing, not 'situation' satisficing ('How Satisficers', p. 42; *The Demands of Consequentialism*, pp. 139–42).

¹² There might be complications here about the welfare levels of dead people, but the example could easily be changed to avoid such problems.

the threshold; and it will never be possible for anyone to get everyone over the threshold. If so, there will be no actual situation in which someone will be permitted not to maximize, and the view turns out to be practically equivalent to maximizing utilitarianism. As Tim Mulgan points out, if the motivation for satisficing is to make morality less demanding, ISALSC is unmotivated.¹³

The more serious problem is that ISALSC still licenses gratuitous harms. Suppose everyone's welfare is well above the threshold level. If ISALSC were true, I could permissibly go around harming people at will for no reason, provided I do not bring anyone's welfare below the threshold level.

There are other problems for situation satisficing. I have not said anything about what counts as the situation after an act is performed, but problems will crop up when we try to do this. For example, should the situation after an act include just states of affairs obtaining shortly after the act? If so, more distant effects will turn out to be irrelevant, which cannot be right. But if the more distant effects are part of the situation after the act, one might be obliged to make up for distant future badness by making things as good as possible in the near future, making satisficing consequentialism practically equivalent to maximizing consequentialism; or, even worse, distant future goodness might relieve one of the obligation to stop massive suffering today. I conclude that no version of the situation view is plausible. What should matter to a satisficer is just what matters to a maximizer: namely, the value of what an act brings about, not the value of a situation.

What can be said about absolute level and comparative satisficing? Absolute level satisficing seems to have one advantage. Maximizers face the problem that sometimes there is no utility-maximizing alternative. For example, take Slote's fairy-tale hero example, and suppose that there is no limit to the value of the reward the hero can request. Maximizing consequentialists seem committed to saying that the hero cannot act rightly in such a situation. This sort of example provides motivation for the view that submaximal acts are permissible. Since comparative satisficers define 'good enough' by appealing to utility-maximizing alternatives, they, like maximizers, are committed to saying that the hero cannot act rightly in situations where there is no utility-maximizing alternative. Absolute level satisficers do not have that problem.¹⁴

¹³ 'How Satisficers', p. 43; 'Slote's Satisficing Consequentialism', p. 127.

¹⁴ Gustaf Arrhenius pointed out to me that absolute level satisficers face this same sort of problem in cases where the utilities of the alternatives asymptotically approach a level at or below the absolute threshold. So the advantage here for absolute level satisficing over comparative satisficing is not very significant.

But both absolute level and comparative versions of SC go wrong with respect to acts that not only fail to cause goodness, but actually prevent goodness. Recall ALSA:

ALSA: There is a number, n , such that: An act is morally right iff either (i) it has a utility of at least n , or (ii) it maximizes utility.

Let n be 20. Suppose that, were I simply to mind my own business and continue sitting on my couch, there would be consequences with intrinsic value of +100. Alternatively, I could get off my couch and undertake a course of action to prevent that outcome. This would involve bringing about a different outcome, with intrinsic value of +20. ALSA entails that both acts are morally permissible. But the act of preventing the better outcome is clearly wrong. In general, it is not permissible to prevent a better outcome by bringing about a worse one.¹⁵ This sort of example applies just as well, *mutatis mutandis*, to CSC1, CSC2, DSC, SALSA and ISALSA.

The sorts of examples Slote takes to provide support for satisficing consequentialism are examples in which someone fails to bring about a better outcome because they bring about an outcome that is 'good enough'. In order for these examples to be persuasive, it is essential that they involve omissions, or allowings. It is permissible, we might think, to forego something better, to allow a better opportunity to pass – to fail to ask for a better reward, or to decline to look for a better room for the stranded motorists. It seems much less plausible to say that it is permissible to *prevent* something better from happening – to intercept someone's reward and replace it with a less good one, or to move the motorists out of the best room and into the merely satisfactory one, for no reason. Steering the world away from a better result towards a less good result should be unacceptable to a consequentialist. But every version of satisficing consequentialism we have looked at so far must allow agents to steer the world towards a worse result in some circumstances.

Of course, we might notice one way in which preventing goodness seems different from allowing goodness to pass: characterizing an act as a prevention makes it seem like a doing, not an allowing. But this distinction is not helpful here. First, consequentialists do not accept a morally relevant distinction between doings and allowings. So for

¹⁵ Compare to Mulgan, 'How Satisficers', p. 42; Mulgan, *The Demands of Consequentialism*, pp. 139–42; and Carlson, *Consequentialism Reconsidered*, pp. 15–17. Mulgan and Carlson's examples also involve gratuitous reductions in utility; but the examples seem perhaps even more convincing when the best consequence can be obtained by simply staying out of the way.

the consequentialist, preventing goodness is on a par with allowing goodness to pass.

Second, even if the consequentialist were to incorporate a distinction between doings and allowings, it would only make things worse. This can be illustrated by considering a principle that might be thought to help the satisficer:

Only Doings Count (ODC): An act of allowing something to happen does not get credit or debit for the value of what is allowed; an act must be a doing in order to get credit for what it brings about.

No consequentialist, satisficing or otherwise, should accept ODC. Consider two situations. In situation one (a Sloten case), I can either (A1) allow a greater good to pass by, thereby allowing a lesser good to exist instead, or (A2) cause the greater good. In situation two (a case like the one I just presented), I can either (A3) prevent the greater good by causing the lesser good, or (A4) allow the greater good to come to pass. ODC entails that A3 is permissible in situation two, since A3 gets credit for the lesser good it brings about, while A4 gets no credit for merely allowing the greater good. And ODC entails that A1 is impermissible in situation one, since A1 gets no credit for merely allowing the lesser good to exist, while A2 would get full credit for bringing about the greater good. These results are exactly backwards. Satisficers of all the varieties discussed so far must believe that doing A1 in situation one might be permissible; and as has been noted, doing A3 in situation two seems impermissible. ODC should be rejected.

It might be thought that the value of what an act prevents should be factored into its utility, perhaps by being subtracted from the value of what it causes, whereas the value of what an act allows should not be factored into its utility. For example, on this proposal, if my act causes states of affairs with intrinsic value of +10, and prevents states of affairs with intrinsic value of -5, the utility of the act would be +15. But this gets the mathematics wrong. Suppose I have two options. One causes states with intrinsic value of +10, the other +5. Doing one prevents the consequences of the other. Then one act has a utility of +5, the other -5. The difference between the utilities is twice as large as it should be. And things obviously get much worse when we consider situations where there are more than two alternatives.

Distinguishing between doings and allowings, or preventions and allowings, will not help the satisficer. If we want to reject the prevention of goodness, we must reject all the versions of satisficing consequentialism discussed so far.

III. SELF/OTHER ASYMMETRY TO THE RESCUE?

There is, however, an entirely different sort of satisficing consequentialism that might seem to be untouched by the sorts of

examples presented in the previous section. The idea is that it is permissible for an agent to forego the best outcome when there would be significant personal sacrifice involved in bringing about the best outcome. In all the counterexamples discussed so far, doing a suboptimal act is completely gratuitous; there is just no reason not to do what is best. In cases where personal sacrifice is required in order to do what is best, there would at least be a reason of some sort to do less than the best. This way of thinking about satisficing is endorsed by John Turri, who says that ‘an outcome *O* is good enough only if *O* is at least as good as the best outcome the agent could have produced in the circumstances without sacrificing something of appreciable personal importance to her’.¹⁶ Call this ‘self-sacrifice satisficing consequentialism’.¹⁷

Of course, since this is only a necessary condition, it does not provide us with an alternative to maximizing consequentialism. Self-sacrifice satisficing consequentialism cannot be evaluated yet, at least as an alternative to maximizing consequentialism, because (i) any act that is right according to maximizing consequentialism is also right according to self-sacrifice satisficing; and (ii) without a sufficient condition for an outcome’s being good enough, there are no wrong acts that self-sacrifice satisficing counts as right. So let us consider a version of self-sacrifice satisficing consequentialism in the spirit of Turri’s proposal that provides both necessary and sufficient conditions for moral rightness:

SSSC: An act, *a*, performed by person *S*, is morally right iff the utility of *a* is at least as great as the utility of any alternative to *a* whose utility for *S* is ‘good enough’.

As we have seen, there are many ways to resolve the vagueness of ‘good enough’, so there will be many versions of self-sacrifice satisficing.

SSSC is not what Slote had in mind by satisficing consequentialism.¹⁸ No personal sacrifice is involved in asking the gods for more good stuff. Finding the best possible room for the stranded motorists would involve only a minimal sacrifice of time on the part of the motel owner; if SSSC counts such a sacrifice as too great, SSSC would turn out to be virtually equivalent to egoism. Of course, it could be argued that Slote was wrong to focus on those examples; perhaps better counterexamples to maximizing involve cases where great personal sacrifice must be made to do what maximizes utility. But those persuaded to be satisficers

¹⁶ Turri, ‘You Can’t Get Away’, p. 2.

¹⁷ Compare to Mulgan’s discussion of ‘cost or effort’ satisficing (*The Demands of Consequentialism*, p. 137).

¹⁸ Slote, *Common-Sense Morality*, p. 47. In light of this fact, Mulgan might justifiably complain that Turri’s objection to his argument is misplaced, since Slote is Mulgan’s target.

by consideration of Slote's examples would have no reason to endorse SSSC.

Mulgan objects to SSSC on the grounds that it takes away the explanatory power of satisficing.¹⁹ Satisficing is introduced to explain why it seems we do not have to perform certain extremely demanding actions. Some versions of satisficing provide an explanation, by entailing that an act can be right just by having consequences of a certain value. SSSC can provide only the following 'explanation': we do not have to perform those demanding actions because they are too demanding.

Mulgan is right to say that there is not much of an explanation here. But it is not clear why this is supposed to be a serious problem. The effort satisficer will just claim that it is a brute fact that morality does not require great sacrifices, and that consequentialism must take this fact into account. Whether that claim is plausible or not depends on what the alternative views are – whether there is a plausible story to be told that explains why great sacrifices are not required. We have already seen that other sorts of satisficing consequentialism are not plausible.

The best objection to SSSC is that it is subject to counterexamples very similar to those that sink other versions of satisficing consequentialism. Suppose that in a certain situation, S can either do nothing, in which case others will benefit greatly, or prevent that great benefit by producing a small benefit, either for himself or for others. Here are S's alternatives:

Alternative	Utility for S	Utility for others	Total utility
A1	0	1000	1000
A2	0	50	50
A3	50	0	50

Suppose that neither A1 nor A2 is good enough for S, but A3 is good enough. Then A3 is permissible according to SSSC. But since A2 is as good as A3, SSSC entails that A2 is also permissible. But A2 cannot possibly be permissible. A1 involves no greater sacrifice than A2, but has much better consequences.

We could tell a little story to explain why S might do A2 in this situation. Suppose S thinks himself unworthy of receiving goods. He also thinks, wrongly, that others are unworthy of receiving goods, though not quite as unworthy as he is. He would rather do A2 than

¹⁹ Mulgan, 'How Satisficers', p. 44; Mulgan, *The Demands of Consequentialism*, pp. 138–9. Mulgan also claims that SSSC is not really a consequentialist theory, since it 'defines moral obligations in terms of costs to the agent' ('How Satisficers', p. 44). But I do not see the force of this objection; it seems fine to call a view consequentialist even if it has some agent-relativity. In any case, this is an argument over a label.

A3, since the benefits go to others rather than himself. And he would rather do A2 than A1, since he thinks others do not deserve the great benefits they would get if he were to do A1. He wants to prevent people from getting those great benefits. So he does A2. Clearly S behaves wrongly here.

So SSSC fares little better than other versions of satisficing consequentialism. Like those other versions, SSSC permits gratuitous prevention of goodness. To be fair, I reiterate that Turri intends to provide only a necessary condition on an outcome's being good enough, and suggests that he would not endorse SSSC.²⁰

A better version of self-sacrifice satisficing consequentialism was suggested to me by Garrett Cullity.²¹ This version attempts to avoid allowing gratuitous prevention of goodness by allowing agents to perform less-than-maximal acts only when doing better would involve a certain amount of self-sacrifice, and when the act is 'good enough' overall. Here is a statement of Cullity's self-sacrifice view, formulated as a version of ALSA:

CSSALSC: There is a number, n , such that: An act, a , performed by agent S, is morally right iff either (i) a has a utility of at least n , and any better alternative is worse for S than a ; or (ii) a maximizes utility.

CSSALSC does not allow for gratuitous prevention of goods. I suspect it (or a comparative-level version of it) might be the most promising version of satisficing consequentialism.

However, this view faces other problems. Consider the following situation. Some money is in an envelope, headed for a charitable organization; if it arrives, it will do a fair amount of good. Jack can intercept that envelope and keep the money for himself. If he does so, less good will result overall, but more good will result for him. Alternatively, Jack can intercept the envelope and give the money to Will. Suppose the utilities work out like this:²²

		Jack	Will	Others	Total
A1	Continue sitting on couch	0	0	+100	+100
A2	Intercept and keep	+20	0	0	+20
A3	Intercept and give to Will	0	+50	0	+50

²⁰ 'You Can't Get Away', p. 2.

²¹ The view was suggested in conversation; Cullity did not endorse it. Something like this view seems to be suggested by Robert Elliot (*Faking Nature*, p. 47).

²² A2 and A3 seem to involve more than one action, but we can avoid this problem easily. Just say that A2 involves pushing one button and A3 involves pushing a different one, redirecting the money different ways.

Now suppose that $n=20$. Since every alternative to A2 is worse for Jack, and A2 exceeds the threshold, A2 is permissible according to CSSALSC. That seems fairly implausible; as before, CCSALSC permits the prevention of a better outcome. Furthermore, since it is not the case that every better alternative to A3 is worse for Jack, A3 is impermissible according to CSSALSC. While it does seem right to say that A3 is impermissible, it is very hard to see how it could be permissible to do A2, but not to do A3.

IV. CONCLUSION

The move to satisficing has been thought to help consequentialists avoid the problem of demandingness. But this is a mistake. The versions of satisficing consequentialism discussed in Section I allow agents to perform suboptimal actions, but are too indiscriminating. As a result, as we saw in Section II, they allow agents to perform such actions even when it is not the least bit demanding to maximize utility. Satisficers might try to avoid this problem by incorporating personal sacrifice into the view explicitly, as in Section III. But in Section IV we saw that such attempts either do not eliminate the problem of gratuitous submaximization, or lead to other similarly unacceptable results. Perhaps more importantly, once we explicitly introduce personal sacrifice into the view, we might wonder what work is left to do for satisficing. In so far as we are worried about demandingness, we might simply adopt maximizing consequentialism with a permission or exception for cases of great personal sacrifice. I conclude that, if satisficing consequentialism is to remain a position worth considering, satisficers must show (i) that the move to satisficing is necessary to solve some problem, whether it be the demandingness problem or some other problem, and (ii) that there is a version of the view that does not permit the gratuitous prevention of goodness.²³

wbradley@syr.edu

Winner of the best paper prize at the Utilitarianism 2005 Conference at Dartmouth College

²³ A version of this article was presented at the Utilitarianism 2005 conference at Dartmouth College; thanks to all those present for their helpful comments, especially Gustaf Arrhenius, Garrett Cullity, Alison Hills, Eric Moore, and Eric Wiland. Thanks also to Jason Kawall, Kris McDaniel, and the students in my Topics in Moral Philosophy course in Fall 2004, especially Mateo Silos Ribas, for helpful discussion.