In Defense of a Version of Satisficing Consequentialism

JASON ROGERS
University of Rochester

In this article, I develop, motivate, and offer a qualified defense of a version of satisficing consequentialism (SC). I develop the view primarily in light of objections to other versions of SC recently posed by Ben Bradley. I motivate the view by showing that it (1) accommodates the intuitions apparently supporting those objections, (2) is supported by certain 'common-sense' moral intuitions about specific cases, and (3) captures the central ideas expressed by satisficing consequentialists in the recent literature. Finally, I offer a qualified defense of the view that consists in showing that it meets Bradley's criteria for being a version of satisficing consequentialism that is 'worth considering'. Specifically, it is a version of SC that solves certain problems for maximizing consequentialism and yet does not permit the gratuitous prevention of goodness.

INTRODUCTION

In a recent (2006) paper, Ben Bradley argues powerfully against satisficing consequentialism (SC). He presents several versions of the view and contends, on the basis of counterexamples purported to show that each of these versions sanctions the gratuitous (or inappropriate[1]) prevention of goodness, that 'every version is unacceptable'.[2] He thus issues a challenge: '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 [or inappropriate[3]] prevention

Sincere thanks go to Earl Conee for his encouragement and helpful input during the course of writing this article. Thanks also go to an audience of faculty and graduate students at the University of Rochester (to whom an earlier version of this article was presented), particularly to Jonathan Matheson and Joshua Spencer.

[1] This qualification will become important for reasons mentioned later in this article. For preliminary remarks on its importance, see n. 3.

[2] Ben Bradley, 'Against Satisficing Consequentialism', *Utilitas* 18.2 (2006), pp. 97–108.

[3] Note that the challenge, as Bradley officially states it here (without my insert), is one only about the *gratuitous* prevention of goodness. Even so, Bradley himself goes on to show that there is at least one version of satisficing consequentialism that does *not* allow the gratuitous prevention of goodness (see Bradley, 'Against Satisficing', p. 107). He still rejects this version, apparently on the grounds that it sanctions an *inappropriate* prevention of goodness. That being so, I have interpreted Bradley's challenge as being that of developing a version of SC that allows neither the gratuitous nor the inappropriate prevention of goodness. The difference between a gratuitous and an inappropriate prevention of goodness is addressed later in this article (see, for example, n. 44). For ease of expression, I will not always include the 'or inappropriate' qualification in what follows.

© Cambridge University Press 2010
Utilitas Vol. 22, No. 2, June 2010
doi:10.1017/S0953820810000099

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

of goodness.'[4] I aim to meet both of these demands in this article. My focus is on meeting condition (ii). But I will also attempt to show that the view that I develop meets condition (i). It resolves what I take to be two serious difficulties for *maximizing* versions of consequentialism: namely, that of not being (intuitively) 'too demanding', and that of allowing for acts that are (intuitively) morally supererogatory. The general problem that the view resolves, then, is that of accommodating certain intuitive data. It thus meets the conditions of Bradley's challenge. I take it that this is not only a necessary, but also a sufficient, condition for rendering the view worthy of consideration.

I begin with some classificatory remarks about the view to be presented. Consequentialism in general, as a view about the moral rightness of acts,[5] is the view that the moral rightness of an act is determined entirely by consequences.[6] Various more specific versions of consequentialism emerge as one specifies whether the determiners of moral rightness are *actual* or *expected* consequences (the resultant views are sometimes called, respectively, 'objective' and 'subjective' versions of consequentialism[7]), whether the relevant consequences are the consequences of the act itself or of (for example) the acceptance of the rule prescribing the act under the circumstances ('act-consequentialism' and 'rule-consequentialism', respectively), whether moral rightness obtains only when the very *best* consequences obtain or when 'good enough' consequences obtain ('maximizing consequentialism' and 'satisficing

[4] Bradley, 'Against Satisficing', p. 108.

[5] In this article, I will assume that 'morally right' and 'morally permissible' are semantically equivalent terms, and so will use them interchangeably. There may be reasons to question this assumption; I make it largely for convenience. Readers with qualms are welcome to substitute 'permissible' for 'right'.

[6] Something very much like this characterization of consequentialism in general is offered by Walter Sinnott-Armstrong, 'Consequentialism', *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (2006) <http://plato.stanford.edu/archives/spr2006/entries/consequentialism/>. The reader is referred to that work for extremely helpful elucidations of various possible versions of consequentialism. As I have presented it here, 'consequentialism in general' also has some affinity with Peter Vallentyne's 'core consequentialism', but the two are not identical (see Peter Vallentyne, 'Against Maximizing Act-Consequentialism', *Contemporary Debates in Moral Theory*, ed. J. Dreier (Malden, MA, 2006), pp. 21–37). Vallentyne's 'core consequentialism' does hold that the rightness (Vallentyne says 'permissibility') of actions is determined entirely by consequences. However, Vallentyne's explication of the two essential theses of 'core consequentialism' appears to rule out rule-consequentialism. Core consequentialism's 'Supervenience' thesis says that 'the permissibility of actions in a given choice situation supervenes on (is fully determined by) the value of *their* consequences' (p. 22, emphasis added), where 'their' apparently refers to the actions themselves. But rule-consequentialism does not maintain that the permissibility of actions is determined by the value of *those actions'* consequences (as Vallentyne himself earlier acknowledges); rather, it maintains that the permissibility of actions is determined by (for example) the value of the consequences of *acceptance of the rule* prescribing those actions in the given situation.

[7] See David Sosa, 'Consequences of Consequentialism', *Mind* 102 (1993), pp. 101–22.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

consequentialism', respectively), and so on. The view defended here is a version of *objective satisficing act-consequentialism*.

Consequentialist theories, including objective satisficing act-consequentialist theories, can vary among themselves over the *feature(s)* of the consequences that determine(s) the moral rightness of acts. The view defended here assumes that the *value* of consequences is what determines the moral rightness of acts, and, more specifically, that it is the *utility* (as a measure of net pleasure minus net pain) of consequences that determines moral rightness.[8] This makes the view a version of *hedonistic evaluative objective satisficing act-consequentialism*.[9]

There remain possible variations on a version of consequentialism with even this level of specificity. Bradley's evaluation of views of this sort, for example, covers at least eight different variations.[10] These variations differ from one another with respect to exactly how the notion of 'good enough consequences' is cashed out. That is, they differ with respect to what fact about the consequences must obtain in order for those consequences to be 'good enough', and for an act to (therefore) be morally right.[11] While I will not investigate each of these variations here, I do wish to begin developing the view to be defended in this article by considering one version in particular.

SITUATIONAL RELATIVITY

Certain versions of SC appeal to the value of a *situation* as the determinant of moral rightness. One such version of SC, developed by Thomas Hurka (1990), says, roughly, that an agent need not maximize

[8] The assumptions are made for two reasons. First, these assumptions make the version of consequentialism being defended here a version that is somewhat standard as far as consequentialisms go. Second, these assumptions seem to be endorsed by the views that Bradley considers. I do not think that anything crucial depends on these assumptions, however, and so the reader is invited to make adjustments in these assumptions should he or she desire.

[9] For more on these categorizations, the reader is again directed to Sinnott-Armstrong, 'Consequentialism'.

[10] See Bradley, 'Against Satisficing'.

[11] For example, one variation, called (by Bradley) 'absolute level satisficing consequentialism' (ALSC), holds that

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.

Another version, called 'comparative satisficing consequentialism 2' (CSC2), maintains that

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.

These views present ways of cashing out what it is for the consequences of an act to be 'good enough': consequences are good enough when they satisfy the non-maximizing utility requirements on the right side of the relevant biconditional. For more on these views, see Bradley, 'Against Satisficing'.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

utility when a situation is already 'reasonably good', i.e. when the value of the situation is already above some absolute threshold. Hurka says, for example, that 'when [the value of] 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 whatsoever to improve it.'[12] The following version of SC, called 'situational absolute level satisficing consequentialism' (SALSC), is a reasonable precisification of Hurka's remarks:

SALSC There is a number, n (the absolute threshold), 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.[13]

The notion, captured by SALSC, that one need not improve a situation when that situation is 'reasonably good' is an instance of the notion that an act is morally right so long as it produces consequences (here, a situation) that are (is) 'good enough', i.e. that an act's being morally right does not (always) require maximizing, but only satisficing.

SALSC is admittedly somewhat imprecise as stated; the notion of a 'situation' is not entirely clear. Bradley suggests the beginnings of an explication: the situation that would obtain after an act is everything that would happen in the world in a very short period of time after the act is performed; further, the value of the situation is a function of the value of the things that happen within it.[14] I will follow Bradley in making the latter claim, though not in making the former. The former claim requires modification. Since any version of consequentialism, as a species of 'consequentialism in general', holds that the moral rightness of an act is determined entirely by *consequences*, the situation that is relevant to an act's moral status is better understood as everything that would happen in the world *as a consequence of the act* in a very short period of time after the act is performed.[15] Fortunately, nothing to be discussed in what follows will turn on this modification.

For, even leaving considerations about the nature of situations aside, it is clear that, as a theory about the necessary and sufficient conditions for an act's being morally right, SALSC simply will not do. Bradley presents the following general counterexample:

Suppose one happens to be in a situation whose value is [already] far above the threshold level [n]. One could [according to SALSC] permissibly kill several

[12] Thomas Hurka, 'Two Kinds of Satisficing', *Philosophical Studies* 59 (1990), pp. 107–11.

[13] Both the statement of the principle and the name of the view itself come from Bradley, 'Against Satisficing'.

[14] See Bradley, 'Against Satisficing', p. 99.

[15] Thanks to John Bennett for the suggestion here.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

people gratuitously, so long as doing so would not cause the value of the situation to drop below the threshold.[16]

Such a result is clearly unacceptable; the completely *gratuitous* killing of several people is, intuitively, never morally right.[17] Further, Bradley is clearly right that SALSC implies that it *is* morally right in the sort of case described. For suppose that n = 100, and that the value of the situation prior to the act is 500.[18] The completely gratuitous killing of several people may produce significant disutility, such that the value of the situation after the act is only 200. Still, the value of the situation after the gratuitous killing is, in this scenario, greater than n. That is, the situation after the act has a value of at least n (100), and so, according to SALSC, the gratuitous killing is morally right. This is the wrong result, and so SALSC is false.

Bradley considers the following modification of SALSC, called 'individualist situational absolute level satisficing consequentialism' (ISALSC), which, he believes, has at least the virtue of not delivering such results:

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.[19]

He maintains that ISALSC avoids delivering the aforementioned results for the following reasons. The counterexample to SALSC just described involves the gratuitous killing of individuals. But (says Bradley) presumably the killing of a person will cause that person's welfare level to drop below n (as long as n is any plausible absolute threshold), and so the act of gratuitous killing will not, according to ISALSC, be morally right.[20] Such gratuitous killings are therefore not endorsed as morally right by ISALSC.

But even assuming that this is right, there are other problems for ISALSC. As Bradley puts it, 'ISALSC still licenses gratuitous

[16] Bradley, 'Against Satisficing', p. 101.

[17] At least, this is true in all cases except perhaps for those (fortunately fantastical) cases where all available alternatives are terrible, and the best of them is an alternative involving the gratuitous killing of several individuals (assuming for the moment that the killing in such a case could even count as 'gratuitous'; the proper understanding of what Bradley means by 'gratuitous' is considered later in this article). This is not the sort of case described here, however; nor is anything like that sort of case considered in this article.

[18] The values assigned here are arbitrarily chosen.

[19] Bradley, 'Against Satisficing', p. 100.

[20] Bradley evaluates ISALSC in this way in 'Against Satisficing', p. 101. He admits that there may be issues regarding the welfare levels of dead people, but notes that the example could be changed to avoid such problems. Presumably, he means that the example could be changed so that it involves causing immense levels of harm to individuals rather than killing them, as in the counterexample that follows in the body of this article.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

harms.'[21] Suppose, for example, that each person's welfare level in a given situation prior to Smith's act is 200. Further, suppose that n = 100. Now let it be that Smith's act is not a gratuitous killing, but is instead a completely gratuitous bashing of Jon's kneecap with a crowbar. After Smith's act, the welfare level of everyone in the situation besides Jon remains at 200; Jon's welfare level, however, drops from 200 to 125. Jon's welfare level therefore remains at or above the threshold, n. In the situation after Smith's act, each person's welfare level is therefore at least n. But then, according to ISALSC, Smith's act of gratuitously bashing Jon's kneecap with a crowbar is morally right. This is false, and so ISALSC is false.

Both versions of SALSC (SALSC and ISALSC) thus face a common problem: they sanction the causing of gratuitous harm. Put in terms more directly pertinent to these theories, the problem is that each one permits an act that gratuitously creates a situation that is worse (either in itself or insofar as things are worse for some individual(s) within that situation) than the situation that obtained prior to the act, so long as the resultant situation has a value still above the absolute threshold. And clearly – that is, intuitively – such a result is unacceptable. To avoid these difficulties, situational satisficing consequentialists (at least of the 'absolute level' sort) must maintain not only that agents are obligated to prevent the situation from falling below some absolute threshold (when possible), but also that agents are obligated not to gratuitously make the situation worse. The theories proposed so far do not deliver this verdict.

That being so, I propose that we allow the intuitions behind Bradley's objections here to push us in the direction of a new theory. The objections considered so far seem to bring out important intuitions about the moral rightness of specific sorts of acts. Acts that cause gratuitous harm are morally impermissible.[22] Our adjustments will attempt to accommodate these new intuitive data, with the aim of developing a theory that fits this evidence from intuitions about certain acts' moral rightness.

An appeal to what I will refer to as 'situational relativity' can be of some assistance here. The general idea behind situational relativity is that there is something morally relevant about the comparative value of the situation that obtains *after* an act is performed relative to the value of the situation that obtains *before* it. A version of SC that appeals to situational relativity can capture the intuitions behind the objections posed by Bradley against SALSC and ISALSC. It can maintain that no morally right act gratuitously makes a situation *worse* overall, even if the situation already has a value above the absolute

[21] Bradley, 'Against Satisficing', p. 102.
[22] See the important qualification in n. 17, however.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

threshold.[23] The following formulation of a version of situational satisficing consequentialism appears to accommodate this data:

SSC (for 'situational satisficing consequentialism') There is a number, n, such that: An act is morally right iff either (i) the value of the situation after the act is at least n, *and* is at least as high as the value of the situation prior to the act, or (ii) the act maximizes utility.[24]

According to SSC, one cannot permissibly make a situation *worse* by causing gratuitous harm or engaging in gratuitous killing. Since both of the counterexamples so far considered present scenarios where the situation *is* gratuitously made worse in such ways, SSC does not succumb to them. It is, in that respect, a better version of satisficing consequentialism.

It is also worth noting that SSC seems to better capture the thoughts of certain satisficing consequentialists, thoughts interpreted by Bradley as in some way supporting either SALSC or ISALSC. For example, Bradley quotes Hurka as maintaining that 'if a state of affairs is already, by absolute standards, reasonably good, [an agent] has no duty whatever to improve it.'[25] Shortly thereafter, he goes on to say (apparently interpreting this remark) that 'Hurka suggests ... that if the situation before an act has at least a certain value, an agent *need not do anything*.'[26] Finding this latter claim implausible (because it allows an agent to permissibly let a situation's value drop *below* the relevant absolute threshold, so long as the *prior* situation has a value above that threshold), Bradley immediately adjusts his interpretation of Hurka's remark to have it that 'what really matters is not the value of the situation *before* the act, but its value *after* the act',[27] and he thus arrives at SALSC. But I do not see that Hurka's remark suggests what Bradley takes it to suggest, and particularly I do not see that it suggests SALSC.[28] Hurka says only that an agent in a situation

[23] Strictly speaking, the theory to be considered does not *quite* say this. I have expressed the claim here colloquially, rather than in the more precise language of the theory. The theory does not actually say anything about a *single* situation's becoming better or worse. Instead, the theory regards the situation before the act as one situation, and the situation after the act as another. It says that the latter situation cannot have less value than the former. Despite this somewhat technical point, I will occasionally express the claims of the theory in a more colloquial manner, for the sake of ease of expression.

[24] I note here that, with the exception of a theory called 'CSSALSC' (to appear later in this article), SSC and all other theories presented hereafter are my own.

[25] Quoted in Bradley, 'Against Satisficing', p. 99. For citation of the original, see n. 12.

[26] Bradley, 'Against Satisficing', p. 100, emphasis added.

[27] Bradley, 'Against Satisficing', p. 100, emphasis in original.

[28] In fairness, Bradley admits that Hurka takes the value of the situation *before* an act to be of significance, and so does not see him as endorsing SALSC outright. My complaints here are just that (i) Bradley appears first to get Hurka's remark wrong when he takes it to suggest generally that an agent in a good-enough situation need not do *anything*, and that (ii) this leads him to modify the view that he takes Hurka to be suggesting in

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

with a value *already* over a certain threshold has no duty whatever to *improve* the situation. This suggests that (i) he explicitly takes the value of the situation *prior to* the act – and not just the value of the situation *after* the act, as in SALSC – to be relevant to determining which acts are morally right, and that (ii) he is at least non-committal regarding an agent's duties (if there be any) with respect to *worsening* (or not worsening) the situation, saying instead only than an agent need not *improve* a situation that is already good enough. Hurka makes exactly these qualifications elsewhere; he writes that 'if the situation is *already* above the absolute threshold, absolute-level satisficing does not require an agent to do anything at all to *improve* it'.[29] He thus seems to be insisting explicitly that the value of the situation prior to the act is relevant, and that situational satisficing consequentialism should only not require that an agent make a situation that is already above some absolute threshold even *better*. He suggests little, if anything, regarding the obligations of agents with respect to the worsening of (or with respect to *allowing* the worsening of) a situation – there is no obvious reason to take him to be suggesting that agents in situations of a high enough value are not obligated to do *anything*. And, in light of just how implausible it is (as Bradley notes) to suggest that an agent in a situation having a value above the absolute threshold can permissibly act to make the situation worse by engaging in gratuitous killing/harming, it seems that the most charitable interpretation of Hurka's remarks is along the very lines suggested by SSC. SSC has it, as Hurka appears to have it, that the values of the situations both before and after the act are relevant to determining moral rightness, and that an agent has no duty to make a reasonably good situation even better – and, importantly, it also has it that an agent cannot permissibly gratuitously make a situation worse.

Unfortunately, in spite of these virtues of SSC, the view is still susceptible to counterexample. Though it does not allow gratuitous harming or killing, it seems to fall prey to an instance of Bradley's more general objection to (almost[30]) all forms of satisficing consequentialism: it sanctions the gratuitous *prevention of goodness* by allowing an agent

the direction of SALSC (and thus away from the view that Hurka actually seems to me to endorse). These issues can be avoided if Hurka's view is instead interpreted along the lines that I suggest (with SSC).

[29] Thomas Hurka, 'Satisficing and Substantive Values', *Satisficing and Maximizing: Moral Theorists on Practical Reason*, ed. Michael Byron (Cambridge, 2004), pp. 71–6. The quotation is from p. 71, emphasis added.

[30] As mentioned in n. 3, Bradley allows that at least one form of SC that he considers does not sanction the *gratuitous* prevention of goodness, but rather sanctions a prevention of goodness that is intuitively inappropriate.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

to gratuitously 'bring about a submaximal outcome in order to prevent a better outcome from obtaining'.[31]

To see this, suppose that n = 100 and that the value of the situation prior to the act that we are about to consider is also 100. Suppose, further, that Smith, in this situation, is just sitting idly on his couch. His daughter skips into the living room with a cookie that she has just baked and sets it on the table in front of him. 'I've baked this cookie for mommy,' she tells him, 'just because I love her. I'll leave it here so she can eat it when she gets home.' Now Smith reflects. He knows that his wife would be extremely delighted by this surprise gift, and he knows that he himself would take some mild enjoyment in seeing his wife so delighted. Still, he would also take some equally mild enjoyment in eating the cookie himself. Further, he knows that his daughter is extremely forgetful, *will* forget (unless reminded) that she ever made such a cookie, and will only be reminded if in fact his wife receives the gift and remarks upon it to her. That said, no one will be harmed if Smith eats the cookie. True, his wife will miss out on some potential delight, but since that delight is not yet hers, Smith will not be actively *harming* her; and likewise for his daughter. So, Smith eats the cookie.

Smith had at least two alternatives before him. He could have eaten the cookie, or he could have simply remained idle and not eaten it. These alternatives and the values of their respective consequences can be represented as follows (let 'Total' = the value of the situation after each act):

Alternative	Utility for Smith	Utility for Daughter	Utility for Wife	Total
Eat	+5	+0	+0	105
Don't Eat	+5	+10	+20	135

As it turns out, the value of the situation prior to the act is 100, which is at least n, and the value of the situation after Smith's act of eating the cookie is 105, which is at least 100. This means that, according to SSC, what Smith did was morally permissible. But it was intuitively *not* morally permissible; it was selfish, and even quite cruel, for Smith to eat the cookie. This is especially so given that he would have experienced no *dis*pleasure in *not* eating the cookie, and could have simply remained on the couch, watching and waiting until a more pleasant overall situation – one which includes the *same* level of mild enjoyment for himself – unfolded before his eyes. SSC has thus

[31] Bradley, 'Against Satisficing', p. 98. Though the objection to follow is an *instance* of Bradley's general objection, the specific counterexample is my own. Bradley himself does not consider SSC (since its presentation in this article is its debut).

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

allowed Smith to actively 'steer the world towards a worse result'[32] than it would have otherwise reached, and for no good reason at all.[33] This is the wrong result, and so SSC is false.

Since we earlier considered an individualistic version of SALSC, it is worth noting that an individualistic version of SSC fares no better with this example. An individualistic version of SSC might be stated as follows:

ISSC There is a number, n, such that: An act is morally right iff either (i) in the situation that obtains after the act, each person's welfare level is at least n, *and* is at least equal to its level in the situation prior to the act, or (ii) the act maximizes utility.

Smith's act in this scenario does not maximize utility. However, in the situation that obtains after his eating of the cookie, each person's welfare level *is* at least equal to its level in the situation prior to the act – again, no one is *harmed* by Smith's action. ISSC therefore delivers the result that his eating of the cookie is morally permissible. So, ISSC is also false.

SELF-SACRIFICE

SSC (and its individualistic variant, ISSC) has allowed Smith to steer the world toward a worse result, and for no good reason at all. Smith has *gratuitously* prevented a better outcome by bringing about an outcome that is less than the best, and SSC has endorsed it. This much is enough to show that SSC is false. But the preceding considerations do not show that it is *always* wrong to prevent a better outcome by bringing about a lesser one. The latter may be permissible when bringing about the better outcome requires enormous personal self-sacrifice.

Thoughts like these are common in the literature on satisficing consequentialism. A view that says something much like this about the permissibility of bringing about sub-optimal outcomes, for instance, is endorsed by John Turri in his defense of satisficing consequentialism. Turri explicitly endorses what he calls the 'Anti-Maximizing Principle (AM)':

AM Some actions are morally permissible despite the fact that an available alternative would promote a better outcome.[34]

Turri apparently endorses AM on the grounds that 'morality will [at times] demand no more than that which can be accomplished without

[32] Bradley, 'Against Satisficing', p. 103.

[33] Though Smith *did* take some pleasure from eating the cookie, presumably this is not really a reason (let alone a good reason) to eat the cookie in this case, since that very same pleasure (plus extra pleasure for others) would have obtained even had he not eaten it.

[34] John Turri, 'You Can't Get Away with Murder *That* Easily: A Response to Tim Mulgan', *International Journal of Philosophical Studies* 13.4 (2005), pp. 489–92. The quotation is from p. 489.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

sacrificing anything of appreciable importance'.[35] On such a view, there are at least possible occasions where the fact that an agent would be required to sacrifice something of 'appreciable importance' in order to bring about the *best* overall outcome makes it permissible for the agent to bring about some suboptimal outcome instead. Insofar as such an agent 'prevents'[36] a better outcome by bringing about an outcome that is less than the best for *this* reason (the thought goes), the agent does not *gratuitously* do so. The view suggested by Turri thus says that on at least some occasions, non-gratuitously (in this sense) preventing a better outcome by bringing about a lesser is morally permissible.

There are, it seems, good intuitive reasons to think that something like this view is right. Suppose, for example, that an agent has two options before him. On the one hand, he can give *everything he owns* to charity. The money and goods which he gives will be put to use to make significantly better the lives of three people in a distant, impoverished land, but he himself will become entirely destitute, live the rest of his life in abject poverty, and suffer serious misfortune related to malnutrition and disease. On the other hand, the agent can give an *extremely large portion* of what he owns to charity, with the result that the lives of *two* distant people are greatly improved, while the third person is now helped only minimally. On this alternative, the agent himself will retain enough to live a reasonably decent – though far from extravagant or what most would call 'well-off' – life. As it turns out, if the agent takes the first alternative, the additional person helped will (unbeknownst to the agent) live a life that is just *slightly* more pleasant than the one that the agent himself will live if he takes the second alternative. This being so, taking the first alternative will

[35] Turri, 'Murder', pp. 489–90.

[36] I use scare-quotes here because there are good reasons for thinking that there are certain situations that meet the description given in the body of this article and yet are not cases of the *prevention* of goodness. There appear to be at least possible situations wherein an agent chooses an alternative that has consequences of lesser value than some other alternative but yet does not *prevent* (according to our ordinary understanding of prevention) the better outcome. Rather, he *precludes* the better outcome by virtue of enacting the alternative with lesser consequences. Consider, for example, an agent choosing between viewing either of two different movies. The agent in fact selects the movie that is less pleasing, and so he enacts the alternative with lesser consequences. The alternative with better consequences is not enacted. Still, it does not seem that the agent *prevents* better consequences from obtaining – it is not as if the better consequences that would result from viewing the other movie are on their way to unfolding (as, for example, they are in the case where Smith eats the cookie) but then are prevented by the agent. Rather, it is just that, by virtue of choosing one alternative, the consequences of choosing the other alternative are precluded. Ultimately, I do not believe that this distinction has any serious repercussions for what is said in the body of this article. It is simply worth mentioning at this juncture. I continue to use the word 'prevent' throughout the article so as to maintain consistency. (The distinction, together with the example of movie-selection, was mentioned to me by Earl Conee.)

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

ultimately result in an outcome that is just *slightly* better overall.[37] But surely the agent's taking of the second alternative, where he still gives *massive* amounts to charity and retains only enough to get by decently, is at least morally *permissible*. It is not the case that the *only* morally permissible option for the agent is that of sacrificing his well-being entirely (undergoing poverty, disease, and malnutrition) in order to make it the case that there is a life that is just *slightly* more pleasant than his would be.[38] While taking that option may be permissible as well, and perhaps even honorable, it is not morally *required*. Thus, some actions are morally permissible despite the fact that an available alternative would promote a better outcome. This is at least so when performing the act with the better outcome requires significant self-sacrifice.

Along similar lines, Peter Vallentyne notes that there is a 'common-sense' notion that some acts are *supererogatory* supporting the idea that it is sometimes morally permissible to perform an act with less than the best consequences. A supererogatory act is one that goes 'above and beyond' what is morally required.[39] Common sense – or, at least, generally shared moral intuition – seems to suggest that there are such acts. Vallentyne notes that, 'according to common sense, devoting all one's income and time to help others is supererogatory'.[40] In cases where such devotion is an available alternative, 'common sense ... holds that some permissible actions are better than others'[41] – that is, it holds that some *non*-best acts are morally permissible. The scenario presented above appears to be one example where this is the case. All of this suggests that the intuitive data here are on the side of a version of satisficing consequentialism that allows enacting non-maximizing alternatives in at least *some* situations. The data are at least not on the side of a version of *maximizing* consequentialism. The data support

[37] I am assuming here that, in the second alternative, the third person remains in a poverty that is equivalent to the poverty which the agent enters in the first alternative, despite the person's being aided to some minimal degree.

[38] Even the infamous 'Singer Solution to World Poverty', a paradigm case of strict moral requirements, does not demand that individuals relinquish *all* of their possessions. Indeed, it explicitly allows that there may be a level of self-sacrifice beyond which no further sacrifice is required. See Peter Singer, 'The Singer Solution to World Poverty', *The New York Times Magazine* (5 Sept. 1999), pp. 60–3.

[39] Here, what is 'morally required' for an agent is the enacting of the most minimally good alternative that the agent could enact while still acting in a morally permissible manner. Enacting at least this minimally good alternative is morally *required* in the sense that the agent fails to act morally (i.e. permissibly) if he fails to do at least this much. Going 'above and beyond' what is morally required thus involves enacting an alternative better than the bare minimum required for the agent's acting in a morally permissible manner.

[40] Vallentyne, 'Against Maximizing', p. 27.

[41] Vallentyne, 'Against Maximizing', p. 27.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

the notion that some morally permissible acts are not (overall) the best, particularly in cases where significant self-sacrifice is involved.

Bradley attempts to develop a version of satisficing consequentialism that accommodates such considerations. The crucial notion to which he appeals, suggested by the above remarks, is the notion that an agent can perform a morally permissible act that does not maximize utility when it is both 'good enough' (has consequences with a utility of at least n) *and* performing the *maximizing* act would be 'worse for' the agent. He calls the version of SC that he develops on the basis of these considerations, suggested to him by Garrett Cullity, 'Cullity's self-sacrifice absolute level satisficing consequentialism' (CSSALSC):[42]

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.[43]

Bradley contends that CSSALSC, despite its initial appearance of promise, falls prey to something like his general objection to all other forms of satisficing consequentialism: it says that it is morally permissible for an agent to prevent some better outcome by bringing about some lesser outcome, even when doing so is (intuitively) clearly wrong. The prevention of goodness that CSSALSC allows is not *gratuitous*, however; it is, instead, a prevention of goodness that is done for a reason that is intuitively just not good *enough*, i.e. for a reason that is at least not sufficient for making the prevention of goodness morally permissible.[44] Call this an *inappropriate* prevention of goodness. Bradley presents an example to this effect as follows:

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

[42] Bradley does not actually mention what the letters in this acronym stand for. I have attempted to discern this from the context.

[43] Bradley, 'Against Satisficing', p. 107.

[44] On this, see Bradley, 'Against Satisficing', pp. 107–8. Bradley says there that CSSALSC 'does not allow for gratuitous prevention of goodness', but adds that it still 'permits the prevention of a better outcome' in cases where such permission 'does not seem right'. This is the version of SC to which I have been alluding in my previous comments about 'gratuitous' versus 'inappropriate' preventions of goodness; see particularly nn. 3 and 30. Bradley's remarks here give some further indication of what it is for an act to be *gratuitous*, in his terms. An act is gratuitous when it is performed for *no* reason (Bradley explicitly says as much on p. 103). In the sort of case presented against CSSALSC, there is *some* reason (for the agent) for not performing the maximizing alternative, but it is apparently not a reason sufficient for making it the case that the agent's not performing the maximizing alternative is permissible. The agent's performance of the non-maximizing alternative is, in my terminology, an *inappropriate* rather than a *gratuitous* prevention of goodness.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

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:

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

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, CSSALSC permits the prevention of a better outcome.[45]

The claim here is essentially this: act A2 is intuitively morally impermissible, and yet CSSALSC says that it *is* permissible, because A2 has a utility of at least n (20), and any better alternative (A1, for example) is such that it is worse for the agent, Jack, than A2. CSSALSC is therefore false.

This conclusion regarding CSSALSC hinges crucially upon what it means for an alternative to be 'worse for' some agent. Alternatives A3 and A1, in the example considered, are apparently worse for Jack, according to Bradley, because they do not make things as good for Jack as some other alternative would. But here is where pressure can once again be applied to Bradley's interpretation of the particular version of satisficing consequentialism under consideration. It seems that there is a more plausible interpretation of what it is for an alternative to be 'worse for' an agent that both better captures the intuitive data considered above *and* makes it the case that CSSALSC does not succumb to this counterexample.

The kind of self-sacrifice that is taken to be relevant in the preceding considerations is the kind that results in an agent's being *less* well off than he was *prior to* the sacrifice. In one of the examples, the sacrifice involves becoming destitute and diseased. In another, the sacrifice requires devoting all of one's income and time to others. Turri's defense of a version of self-sacrifice satisficing consequentialism is motivated in part by an example where an agent, in order to perform the maximizing act, must go from being a healthy human being to being entirely paralyzed.[46] All of this suggests that Bradley's interpretation of the phrase 'worse for' in CSSALSC fails to result in a theory that best fits the intuitive data or that best fits what self-sacrifice satisficers appear to have in mind. Rather, the general idea seems to be that an alternative

[45] Bradley, 'Against Satisficing', pp. 107–8.
[46] See Turri, 'Murder', p. 491.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

is 'worse for' an agent if and only if that alternative makes the agent *less* well-off than he was prior to acting. An alternative is worse for an agent if and only if the alternative has consequences with net disutility for the agent relative to the agent's prior level of welfare – in this case, things are worse for the agent after enacting the alternative than they were prior to doing so. In short, saying that an alternative is 'worse for' some agent really amounts to saying that it *makes things worse for him*.

This interpretation of 'worse for' can be explicitly incorporated into a version of satisficing consequentialism. Consider CSSALSC2:

CSSALSC2 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 to a is such that, were that alternative to be enacted by S, S's resultant personal welfare level after that enaction would be less than it was prior to that enaction; or (ii) a maximizes utility.

CSSALSC2 does not succumb to the counterexample to CSSALSC. A2 has a utility of at least n (20), but it is not the case that any better alternative is such that, were that alternative to be enacted instead of A2, the agent's resultant personal welfare level after the enaction would be *less* than it was prior to the enaction. Indeed, of alternatives A1, A2, and A3, CSSALSC2 says that only A1 is morally permissible. This is exactly in line with our intuitions regarding the case. This is reason to think that CSSALSC2 is a more plausible view than CSSALSC.

Even so, CSSALSC2 is still not fit to accommodate all of the intuitive data. Suppose that in alternative A1, Jack's foot falls asleep as he continues sitting on the couch, and so enacting A1 results in −1 utility for Jack. With the utilities thus modified, CSSALSC2 delivers the result that it is permissible for Jack to intercept the envelope on its way to charity and give it to Will. This still does not seem right.[47]

The reason it does not seem right, I think, is that the sacrifice is not significant enough. The sacrifices considered above are not sacrifices that involve some small amount of disutility. Rather, they are sacrifices that involve things becoming *significantly* worse off for the agent. *These* are the sorts of sacrifices that make it the case that it is permissible for an agent to bring about some lesser alternative. Turri, in a previous citation, explicitly uses the language of 'appreciable importance', saying that sacrifices requiring the agent to sacrifice something of such

[47] It is true that CSSALSC2 delivers the verdict that A1 is *also* morally permissible in this case (since it maximizes utility). The problem is just that, of the three alternatives, it seems intuitively that A1 is the *only* permissible one; A3 is not also permissible. But CSSALSC2 says that A3 *is* permissible, since it meets CSSALSC2's condition (i): A3 has a value of at least n, and the only overall better alternative, A1, is such that its enaction results in a decrease in Jack's personal welfare.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

importance are the ones that (can) make it the case that it is permissible for the agent to bring about consequences that are less than the best.

These points can also be incorporated into a version of satisficing consequentialism, though something more should first be said about what it is for a sacrifice to be 'significant' or 'of appreciable importance'. What I will say is the following. It seems to me that these notions involve a certain amount of 'relativity', as the following example shows. Suppose that an agent is faced with the alternatives of either wrapping the gift which he has purchased for his wife or not doing so.[48] Suppose that if the agent wraps the gift, his wife will be quite pleased with it, but he will endure a stinging paper cut (and not enough pleasure to 'make up for it' on his end). Suppose that if he does *not* wrap the gift, his wife will still be rather highly pleased with it, and he will endure no paper cut. Even though the former alternative will (suppose) result in *slightly* higher overall utility, it seems that the latter option is at least morally *permissible*, given that the agent is, in that alternative, able to avoid suffering the paper cut, and his wife is still highly pleased. Perhaps this means that, given the only *minor* difference in resultant overall utilities between the two alternatives (the agent's wife is only *slightly* happier in one as compared with the other), the sacrifice of enduring the paper cut is, *relative to* that very minor difference, significant enough to be relevant to moral permissibility. This suggests, perhaps, that the level of significance of a sacrifice required for the sacrifice to be morally relevant (i.e. of 'appreciable importance') is determined by considering the disutility of the sacrifice for the agent relative to the difference in overall value between the available alternatives. If the sacrifice is very small relative to the difference in overall utility, then it will not be of appreciable importance. But if, relative to the difference in overall utility, the sacrifice is not comparatively minor, it *will* be morally relevant. I will call a sacrifice that is significant in this latter sense a *comparatively significant* sacrifice.

Here is the resultant view:

CSSALSC3 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 to a is such that, were the alternative to be enacted by S, S's resultant personal welfare level after that enaction would be *comparatively significantly* less than it was prior to that enaction (that is, the decline in S's personal welfare level, relative to the difference in overall utilities between act a and the overall better alternative under consideration, is sufficient to make for comparatively significant sacrifice); or (ii) a maximizes utility.

[48] Something generally like this example is presented in Vallentyne, 'Against Maximizing', though toward different ends.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

This view appears to avoid the counterexamples to CSSALSC presented by Bradley and myself, and to accommodate all of the aforementioned intuitions about self-sacrifice. Returning to the modified example of Jack and the letter to charity, for instance (i.e. the version of this example wherein Jack's foot falls asleep if he takes alternative A1), CSSALSC3 delivers the verdict that only option A1 is morally permissible. This is because neither of the non-maximizing alternatives, A2 or A3, is such that any better alternative (namely, A1) involves a *comparatively significant* self-sacrifice for Jack. Relative to the overall differences in utility between the better alternative, A1, and those alternatives not involving any sacrifice on Jack's part – a difference of 79 in the case of A2 and a difference of 49 in the case of A3 – Jack's sacrifice with a disutility of −1 is not comparatively significant.[49] In Turri's language, it is not of 'appreciable importance'. The verdict that CSSALSC3 delivers here is exactly in accord with our intuitions about the case.[50] As such, CSSALSC3 appears to have substantial promise.

A FINAL, HYBRID VIEW

CSSALSC3 is a version of satisficing consequentialism that avoids Bradley-style objections. Showing this is already showing something significant; it is showing that there is at least one version of satisficing consequentialism that meets his condition (ii), that of not allowing the gratuitous (or inappropriate) prevention of goodness. We have also seen some reason, mentioned above, to think that CSSALSC3 accommodates the intuitive data suggesting that some acts are supererogatory, which intuitive data versions of *maximizing* consequentialism appear to have difficulty accommodating. This being so, CSSALSC3 is a version of SC that meets the whole of Bradley's challenge. But I wish ultimately to arrive at what I think is the *most* plausible version of satisficing consequentialism *vis-à-vis* maximizing consequentialism, and I do not think that CSSALSC3 is it. It succumbs to an objection unlike those considered so far.

The objection is this: despite the fact that CSSALSC3 is a satisficing view, it is still 'too demanding'.[51] More specifically, the view asks too

[49] Such judgments will obviously be largely intuitive. I do not claim to have provided a precise algorithm for determining comparative significance.

[50] The reader is encouraged to apply CSSALSC3 to the other cases considered, keeping in mind that CSSALSC3 allows satisficing *only* when maximizing requires the agent to make a sacrifice that makes things *worse* for the agent to a degree that is comparatively significant.

[51] I use scare-quotes here because the claim that a view is too demanding is now rather well known in the consequentialist literature. Bradley ('Against Satisficing') and Vallentyne ('Against Maximizing') cite some of the relevant works on demandingness,

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

much in certain cases where the situation is already 'good enough'. Suppose, for example,[52] that Trent has just eaten a quite respectable lunch and finds himself no longer hungry, though also not completely stuffed. If Trent were to eat a candy bar, he would enjoy it. Further, there is (suppose) just such a candy bar at the snack stand nearby. It is available for free, as part of a promotion running in Trent's area. In this situation, Trent has at least two alternatives available, one being that of remaining on the park bench where he has just eaten lunch, and the other being that of walking to the snack stand, getting the candy bar, eating it, and then returning to the bench. The latter alternative, suppose, is the alternative that maximizes utility. Taking that alternative requires no sacrifice on Trent's part; things will not become *worse* for Trent should he consume the candy bar. Taking the other alternative, though, is still pretty good; it results, suppose, in a situation with a value greater than the relevant threshold. However, *because* taking the maximizing alternative requires no comparatively significant sacrifice on Trent's part, CSSALSC3 delivers the result that taking this lesser alternative is *not morally permissible*.[53] And that does not seem right.[54] Surely it is *permissible* for Trent to relax on the park bench rather than eat a candy bar, even if it is not *best*, and even if the best requires no significant self-sacrifice. This seems to be a kind of satisficing that is an intuitively morally permissible alternative to maximizing.[55]

We saw earlier that SSC did a decent job of accommodating these latter intuitions. Its flaw was that it allowed the gratuitous prevention of the best results. It was for precisely this reason that we moved to a consideration of views which allowed the prevention of best results only when significant self-sacrifice would be required to achieve them,

most (if not all) of them focused on the problem of demandingness for versions of maximizing consequentialism. This is worth noting because, if the view developed here *avoids* the demandingness problem while maintaining the other virtues of CSSALSC3, then it satisfies condition (i) of Bradley's challenge in at least two respects, by virtue of solving two problems facing maximizing consequentialism: the demandingness problem and the (related) problem of accommodating the intuition that there are morally supererogatory acts.

[52] This example is inspired by suggestive remarks toward a similar case in Hurka's 'Satisficing and Substantive Values'.

[53] Note that at least hedonistic evaluative objective maximizing act-consequentialist theories (i.e., traditional act-utilitarian theories) will deliver this same result. In a nutshell, this is the traditional demandingness problem.

[54] Commenting on a similar scenario, Hurka remarks, 'Surely that is implausible . . . If a situation is already reasonably good . . . there is no moral demand to make it better.' See Hurka, 'Satisficing and Substantive Values', pp. 72–3.

[55] It is worth noting here that, as I also hint in n. 51, the intuition driving the demandingness problem is, in some sense, 'another side' of the intuition that some acts are morally supererogatory. It is at least similar in that it suggests that there are morally permissible non-best acts.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

and arrived at CSSALSC3. This suggests that SSC and CSSALSC3 each face a particular difficulty that the other avoids. CSSALSC3 is too demanding; SSC is not. SSC allows the gratuitous prevention of the best results; CSSALSC3 does not. The situation is such that not attempting a combination of at least some elements of the two views (in order to avoid both problems) would be negligent.

Consider the following attempt, which I will call 'situational absolute-level self-sacrifice satisficing consequentialism', or SASSC. SASSC incorporates, from CSSALSC3, the notion that comparatively significant self-sacrifice is morally relevant, in the sense that it can play a role in making it the case that some act with consequences less than the best is morally permissible. It also incorporates, from SSC, the notion that the value of the situation after an act, relative to the value of some other situation(s),[56] can play a role in making it the case that not maximizing is morally permissible.

SASSC There is a number, n, such that: An act, A, performed by agent S, is morally right iff either (i) the value of the situation after A is at least n, and is at least as high as the value of the situation prior to A, *and* any overall better alternative to A,[57] A∗, is such that: [were A∗ to be enacted instead of A, *either* S's resultant personal welfare level after the enaction of A∗ would be comparatively significantly less than it was prior to the enaction of A∗, or the value of the situation after the enaction of A∗ would not be appreciably greater than the value of the situation after the enaction of A]; or (ii) A maximizes utility.[58]

SASSC is our most complex version of satisficing consequentialism yet. Some of its elements require explanation. I will provide these explanations and then go on to argue that the theory avoids all of the problems that the previous views faced.

It is first worth stating SASSC in somewhat more colloquial terms, so as to better understand exactly what it claims. SASSC says that there is some absolute threshold, n, that the value of a given situation must be at or above in order for satisficing to be permissible; otherwise, maximizing is required. When the value of the situation *is* above the threshold, satisficing is permissible only if a number of additional conditions are met. First, it must be that the satisficing action results in a situation with a value at least as high as the value of the situation preceding the action; no permissible satisficing action can make a situation worse.[59] Second, it must be that any *alternative*

[56] Here, as SASSC indicates (and as will be explained in what follows), both the value of the situation prior to the act *and* the value(s) of the situation(s) that would result from enacting other available alternatives are relevant.

[57] Better alternatives are just those that produce situations with higher overall values.

[58] Brackets and emphases are included here only as an aid to the reader.

[59] Many readers, particularly those inclined toward satisficing consequentialism, may worry that this element of SASSC is still too demanding. Such readers are encouraged to

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

action that results in greater overall utility than the satisficing action satisfies at least one of the following conditions: either (i) enacting the alternative instead of the satisficing action would require a comparatively significant self-sacrifice (i.e. a sacrifice of 'appreciable importance') for the agent, or (ii) the overall utility that would result from the enaction of the alternative is not *appreciably greater than* the value of the situation that results from the enaction of the satisficing alternative. When all of the preceding conditions are satisfied, SASSC says that the act is morally permissible.

Stating SASSC in both its formal and colloquial terms has required appeal to the (new) notion of a situation's value being 'appreciably greater' than the value of some other situation. This needs explanation. As I am thinking of things here, whether or not the value of one situation is *appreciably greater* than the value of another is determined by comparing the difference in value between the two situations with the *cost to the agent* of bringing about the situation with the greater value as opposed to the one with the lesser.[60] If, for example, the difference between the values of the situations is quite minimal, and bringing about the one with the greater value has a large cost to (i.e. requires a large amount of exertion or energy from) the agent, then the value of the better situation is *not* appreciably greater than the value of the lesser. If one situation is somewhat greater in value than another and bringing it about exacts quite *minimal* cost, however, then the value of the situation is indeed appreciably greater. Here, it is important to note that the 'cost to the agent' is not to be understood in terms of sacrifice (as a *loss of personal welfare*) for the agent. Rather, 'cost to the agent' is intended as a measure of the amount of energy, work, or exertion required on the agent's part to bring about some alternative. Such work or exertion need not involve a sacrificial loss of personal welfare on the agent's part.

Given this understanding of one situation's value being 'appreciably greater' than the value of another, SASSC is not too demanding in the way previously considered. It does not say that it is morally

'stick it out' for the moment. Due to organizational concerns, I address this worry later, in the article's brief appendix.

[60] Questions might arise here: why have I now mentioned 'cost to the agent' instead of simply appealing again to the *sacrifice* required for the agent to bring about one situation as opposed to another? And how (if at all) are costs to the agent and sacrifices related? Each of these questions is answered at the end of the paragraph in the body of this article. But put briefly, 'cost to the agent' is intended as a mere measure of the energy, work, or exertion required for the agent to enact an alternative; it need not be that a cost to the agent is (all things considered) a *sacrifice* on the agent's part, resulting in a loss of personal welfare for the agent.

In certain ways, appreciable greatness exhibits a 'relativity' similar to that exhibited by 'comparative significance' in my earlier discussion of comparatively significant sacrifices.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

impermissible for Trent to relax on the bench rather than eat a candy bar. At least, it does not say this given certain assumptions. Assuming that the value of the situation in which Trent eats the candy bar is not appreciably greater than the situation in which he remains sitting on the bench – which seems plausible: the transition to the former situation exacts some cost (Trent's leaving his bench, walking to the snack bar, requesting the chocolate, etc., when he feels no real desire to do so) in exchange for *extremely minimal* value – then it is morally permissible for Trent to remain sitting on the bench.[61] The alternative in which Trent eats the candy bar instantiates a situation with a value not appreciably greater than the value of the situation that obtains when Trent remains on the bench. In summary, the value of the situation preceding Trent's act of remaining on the bench is, we are supposing, above the absolute threshold (it is 'good enough'), his act does not make the situation any worse, and the overall better alternative, were it to be enacted, would not result in a situation of appreciably greater value than the one that results from the act under consideration. So, according to SASSC, Trent's act of remaining on the bench is morally permissible.

SASSC also does not fail in the opposite direction, by allowing the inappropriate or gratuitous prevention of goodness. Consider again the case of Jack and the letter to charity, for example. Alternative A1 comes out as morally permissible, according to SASSC, and alternatives A2

[61] I admit that some may feel inclined to deny the assumption here, particularly given that I have indicated that a situation having only somewhat greater value than another, and the actualization of which exacts quite minimal cost to the agent, *is* appreciably greater than another (see the previous paragraph in the body of this article). I thank Earl Conee for pointing this out to me. In response, I have two rejoinders. First, I do not think that it is right to deny the assumption here, because, in the situation that I am envisioning, the enjoyment of the candy bar by Trent results in a really *quite* minimal and qualitatively superficial additional value. With respect to *this* value, the costs to Trent of putting aside his relaxation on the bench, walking to the snack stand, requesting the candy bar, etc., seem significant enough. Thus, I think it is right to say that the cost to Trent here is such that the value of the resulting situation is not appreciably greater than the value of the situation wherein Trent remains on the bench. But second, leaving these considerations aside, I do not see that the fact that someone may deny the initial assumption is a problem for SASSC. To the extent that someone thinks that one situation really *is* appreciably greater than another – and this is what is involved in denying the assumption here – they ought also to think that the correct moral theory will require someone to bring about that appreciably greater situation (at least, barring significant self-sacrifice), and so the fact that SASSC will, if the assumption is false, say that Trent's relaxing on the bench is impermissible is not bad news for the theory. It does not seem too demanding for a theory to require someone to bring about a better situation when that situation is appreciably greater than an alternative and when doing so requires no significant self-sacrifice. Turri makes a similar claim in a related context, remarking, 'As far as I can tell, there is no interesting sense in which it is *too demanding* to require someone to promote a better outcome so long as it doesn't require any [appreciable] personal sacrifice on their part' (Turri, 'Murder', p. 491, emphasis in original).

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

and A3 come out as *im*permissible, which is just the right result. This is because alternatives A2 and A3 have a better alternative, A1, which is such that (i) its enaction does not require significant self-sacrifice and (ii) its value *is* appreciably greater than the value of the situations that obtain in the other alternatives (A2 and A3). This shows that SASSC avoids the problems that remained for CSSALSC3 and for SSC. That is an advantage of the view.

The view's only obvious disadvantage – one that I have already noted – is its complexity. But given the theory's demonstrable merits, I do not regard this as a serious disadvantage. In fact, I do not regard it as a disadvantage at all; given the apparent complexity of moral matters, the theory does not seem more complex than should be expected. This being the case, I recommend it for further consideration. SASSC is a version of satisficing consequentialism that (i) does not allow for the gratuitous or inappropriate prevention of goodness and that (ii) solves the problem of accommodating the intuitive data regarding both supererogation and demandingness that other versions of consequentialism fail to accommodate. It thus meets Bradley's challenge and stands as a version of satisficing consequentialism worth considering.

APPENDIX: PREEMPTING A POTENTIAL WORRY ABOUT SASSC

There may be some worry that SASSC remains too demanding.[62] Suppose that an agent finds himself in a situation with only two alternatives available to him, one of which very minimally increases overall utility but requires drastically significant self-sacrifice, and the other of which requires no such self-sacrifice but very minimally *lowers* overall utility. It seems plausible to think that it could be morally permissible for an agent to take the second of these options,[63] given the small difference in resultant overall utility and the drastic nature of the sacrifice required for maximizing. But SASSC does not allow this. SASSC entails that no permissible non-maximizing action can make a situation worse, even *minimally*. So, if SASSC is right, then the only morally permissible option for the agent is to make the drastic sacrifice and minimally increase overall utility.

There are, I think, at least two ways of eliminating this apparent difficulty for SASSC. The first accepts the claim that the second of the aforementioned options is morally permissible, but modifies SASSC so as to make it say that no permissible non-maximizing action makes a

[62] See n. 59 in this regard. I thank Jonathan Matheson for pressing me on this point.
[63] At least perhaps so long as the value of the resulting situation does not fall below the bare minimum marked by the absolute threshold.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

situation *appreciably* worse.[64] The idea here is that an agent makes a situation appreciably worse only if the difference in value between the resultant (worse) situation and the prior situation is large compared with the cost to the agent of bringing about any situation *better than* the resultant situation. The resultant view, SASSC2, is SASSC with this modification made explicit:

SASSC2 There is a number, n, such that: An act, A, performed by agent S, is morally right iff either (i) the value of the situation after A is at least n, and the situation after A is not *appreciably worse* than the situation prior to A, and any overall better alternative to A, A∗, is such that: [were A∗ to be enacted instead of A, either S's resultant personal welfare level after the enaction of A∗ would be comparatively significantly less than it was prior to the enaction of A∗, or the value of the situation after the enaction of A∗ would not be appreciably greater than the value of the situation after the enaction of A]; or (ii) A maximizes utility.

Since the lesser situation in the aforementioned scenario is *not* appreciably worse than the prior situation – the difference between the lesser situation and the prior situation is small compared with the cost to the agent of bringing about the better alternative – SASSC2 says that the non-maximizing act is morally permissible.[65]

The second way of responding keeps SASSC as it is and denies that the case under consideration is a counterexample by denying the claim that the non-maximizing act is morally permissible. While the supposition that it *is* morally permissible is initially plausible, it does not obviously continue to be so when the case is considered with more explicit care. It is important to note, for instance, that the agent's own personal welfare level factors into the calculation of the overall utility in the resultant situation. Thus, given that the agent makes a *drastic* sacrifice in bringing about the maximizing alternative in the case under consideration (and so has his personal welfare level substantially decreased), and yet the maximizing alternative still makes things *better* overall,[66] it must be that the agent's substantial sacrifice is counterbalanced by a substantial gain of utility by others (either individually or aggregately). Once this is noticed, the claim that SASSC is too demanding in requiring the agent to enact this alternative has less plausibility. Sometimes – plausibly, when the result is a substantial increase in overall utility – morality appears to require self-sacrifice. So far as I can tell, common sense endorses this claim.

[64] Earl Conee suggested this modification to me in conversation.

[65] If the situation *were* appreciably worse, then SASSC2 would require maximizing, but that is not obviously incompatible with our intuitive judgments about the case, given the appreciable difference between the outcomes.

[66] The point made here would apply even if the maximizing alternative were only to keep the situation *equally* as valuable as before, however.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press

Thus, the fact that SASSC says that sacrifice is required in these sorts of cases does not clearly show that SASSC is too demanding, and so is no objection to the view. SASSC still allows that self-sacrifice can make satisficing morally permissible in certain cases (e.g. when there are two or more alternatives that make a situation better or equal, at least one of which does not require significant self-sacrifice), but it also allows that morality at times *will* require significant self-sacrifice (e.g. when one can make the situation better by sacrificing, and otherwise can only make it worse). To me (and, I hope, the reader), this feature of SASSC seems just right.[67]

jroger12@mail.rochester.edu

[67] Indeed, because of this and its other features, SASSC will be able to avoid the charge, leveled by Timothy Mulgan (see Timothy Mulgan, 'How Satisficers Get Away with Murder', *International Journal of Philosophical Studies* 9 (2001), pp. 41–6), that satisficing views allow agents to 'get away with murder'. A lengthy consideration of this issue is beyond the scope of this article. Briefly, however, Mulgan's objection is based on an example in which an individual, Mary, finds herself standing on a bridge next to a heavy sandbag, a light sandbag, and an innocent bystander, Bob, while overlooking a trolley line on which there is a trolley carrying ten passengers that is hurtling out of control toward a nearby cliff. In the example, if Mary does not do something, it is certain that the trolley will plunge over the cliff and the ten passengers on board will die. Mulgan worries that satisficing views will allow Mary to permissibly shoot Bob, causing his body to fall and stop the trolley, instead of simply throwing the heavy sandbag onto the tracks and stopping the trolley thereby (even though neither action requires sacrifice on Mary's part, or even if the latter requires very little sacrifice). Fortunately, SASSC does not allow this.

Since I have mentioned Mulgan, it is worth also saying a word about a general concern that he raises for satisficing views that, like SASSC, appeal to sacrifices and 'costs to agents'. In response to Turri's particular view, Mulgan (see Timothy Mulgan, 'Reply to John Turri', *International Journal of Philosophical Studies* 13.4 (2005), pp. 493–6) contends that such views are not really versions of consequentialism, since they 'define moral obligations in terms of costs to the agents, not just in terms of features of outcomes' (Mulgan, 'Reply', p. 494). But I do not see the force of this objection. That an agent becomes such that he has had to endure a cost of some value x by bringing about some alternative just *is* an outcome of the act of bringing about that alternative, so far as I can tell. Additionally, there are non-hedonistic versions of consequentialism that can appeal to other features of outcomes as morally relevant besides net pleasure and pain, such as, for example, the fact that the outcome is such that there is an agent who had to undergo significant self-sacrifice to bring it about (for more on species of consequentialism with different value theories, see Sosa, 'Consequences', and Sinnott-Armstrong, 'Consequentialism'). Finally, there remains the fact that, Mulgan's concern aside, self-sacrifice satisficing views constitute a subset of the views against which Bradley objects, and so whether such views are indisputably consequentialist is an ultimately tangential concern – this article remains a defense of the sort of views against which Bradley objects using the moniker 'satisficing consequentialism'.

https://doi.org/10.1017/S0953820810000099 Published online by Cambridge University Press