

Is Satisficing Necessary or Just 'Good Enough'?

One of the major criticisms of traditional act consequentialism (TAC), as compared to common sense morality, is that it does not provide an agent enough options. An agent must always perform an act that has at least as much utility as any other. That is to say, she must perform the optimal act, or one of the optimal acts if the ranking for optimal act produces a tie. This has motivated consequentialists to produce a theory that is consistent with common sense morality on this point, but still consequentialist. One such theory put forward by Michael Slote is satisficing consequentialism (SC). However, as some have shown, SC produces some counterintuitive implications, such as the gratuitous prevention of goodness, which make it seem implausible. But I will show why this is, and compare it to an option that is considered more plausible, namely Schefflerian Utilitarianism¹, which allows for egoistically-adjusted utility. Then I will demonstrate why SC is still a viable option for the self/other problem, viz. the problem that common sense morality seems to show that there is a self/other asymmetry, by showing how SC can be formulated to get around the gratuitous prevention of goodness objection.

First I will lay out the motivation for SC and then explicate the theory. Next I will show how most versions of SC allow for the gratuitous prevention of goodness. Third, I will expound on the exception, that is, the version of SC that does not allow for gratuitous prevention of goodness, namely self-sacrificing satisficing consequentialism (SSSC). Fourth, I will explore whether this theory needs to be satisficing and show how it can be explained without the satisficing element. After that I will contrast SSSC with a non-satisficing self-sacrificing consequentialist view. And lastly, I will show why the

satisficing view might be desirable considering some of the objections that can be levied against other self/other consequentialist views.

I

To many, TAC seems to be too restrictive and too demanding. This is because TAC does not give an agent options other than those acts that are optimally ranked, and it requires that the agent perform the optimal act. That is, the agent must always perform an act that is ranked at least as high as any other act. However, as Michael Slote points out

Could not someone who held that rightness depends solely on how good an act's consequences were also want to hold that less than the best was sometimes good enough, in other words, that an act might qualify as morally right through having good enough consequences, even though better consequences could have been produced in the circumstances?²

Here Slote is alluding to a feature of common sense morality, i.e. that an agent ought to have options in most cases. Furthermore, he shows how these types of cases are common in our everyday, non-moral reasoning by appealing to an economic example of satisficing.

An example of satisficing is making a decision of where to eat for dinner. Spouses may start running through options until they find one that is good enough without identifying the optimal dinner location. Furthermore, they may later think of a better option, but decide to reject it for the lesser option *ceteris paribus*. This doesn't seem bad. In the same way, it doesn't seem bad for an agent to select a moral option available to him that is not the optimal option, but rather is simply good enough.

An example of moral satisficing given by Slote is that of the heroine who is rewarded with one unconditional wish³. The heroine chooses as her reward enough

money for her and her family to live comfortably for the rest of their lives. We don't, Slote says, find the heroine irrational for choosing less than she could even if choosing more would make her and her family better off. This type of 'good enough' act should be captured by morality, but TAC would consider the act impermissible. This case and related non-fantastical cases motivate the move toward a moral theory that captures this intuition, namely SC.

SC can be formulated different ways, but two main categorizations are absolute level (ALSC) and relative satisficing consequentialism (RSC). ALSC says that there is some absolute level that an act must rank at in order to be considered good. For instance, if the absolute level of goodness is 100 hedons, then any act that is 100 hedons or above is permissible. Any act below 100 hedons is impermissible. RSC, on the other hand, assigns a relative level by way of a percentage that an act must rank at compared to the optimal act. For example, if the optimal act ranked at 240 hedons and the relative level was 90% then any act below 216 hedons would be impermissible. Or if the optimal act ranked at 400 hedons, then any act below 360 hedons would be impermissible.

To see how this works consider the following example where an agent decides to make cookies for a homeless shelter. He has enough ingredients to make two cookies for each homeless person and the ingredients will go bad tonight if he doesn't use them. Further suppose his oven is big enough to cook all the cookies at once so there is no extra effort for him to make enough for each homeless person to have two. Also, assume that each homeless person would be satisfied with one cookie although two would make them a little happier. If the agent makes the smaller batch, is that wrong? Not doing anything

for the homeless seems wrong given the option, but providing one cookie each instead of two certainly seems to be an act worthy of being considered good. Certainly it's not optimal, but it seems good enough. This example might produce the following rankings:

	U(x)	TAC	RSC	ALSC
Do nothing	0	Impermissible	Impermissible	Impermissible
Cook smaller batch	100	Impermissible	Permissible	Permissible
Cook larger batch	110	Permissible	Supererogatory	Supererogatory

I take it that the major motivation for SC is to provide an agent options that maximizing forms of consequentialism such as TAC do not provide, and as Slote states, to bring consequentialism more in line with common sense morality⁴. To do this SC needs to avoid any major deviations from common sense morality or be able to explain them. In the next section I focus on some of the objections to SC and whether these lead to counterintuitive implications.

II

There are many variations of ALSC and RSC, but most of them have a similar problem. The problem is that they allow the gratuitous prevention of goodness. To see why consider the following case. Imagine an agent's parent is about to write a check for \$9,000 to benefit Oxfam, and suppose the agent's parent is sufficiently rich such that the amount will have minimal impact on the parent and none on the agent. Now suppose the agent could ask the parent to give \$900 of that money to the agent's fraternity which will use it frivolously and hence provide no extra utility, leaving only \$8,100 Oxfam. If we assume RSC is true and further assume a level of 90 percent, then writing \$8,100, or \$9,000 to Oxfam are both permissible acts. However, it seems like the only right thing to do is to allow the parent to write the check for \$9,000. For suppose it costs \$300 to

provide one years supply of food. Requesting the \$900 is allowing the gratuitous prevention of goodness, i.e. three people fed for a year. ALSC doesn't fare any better. Imagine that the absolute level is 810 and that giving \$1,000 to Oxfam ranks at 100 on a linear scale up to giving \$9,000, which ranks at 900. This provides the same implications as RSC. Requesting the \$900 would rank at 810 and therefore be permissible as would allowing the giving of \$9,000. Oddly though, supposing you could request and receive \$910 instead, this seems like it would be impermissible rather arbitrarily since it would only have a value of 809.

Oxfam	U(x)	ALSC	RSC	TAC
\$8,090	809	Impermissible	Impermissible	Impermissible
\$8,100	810	Permissible	Permissible	Impermissible
\$9,000	830	Permissible	Permissible	Permissible

Whereas causing the giving of only \$8,100 is impermissible on TAC, it's not arbitrary. It's impermissible because there is a better act that can be performed. However, SC has set an arbitrary lower limit to how much utility one can permissibly perform. So it's not clear why \$8,090 is impermissible on SC. This objection applies to the case mentioned in the previous section pertaining to the homeless shelter and baking cookies. I think this objection brings to light why it seems wrong to refrain from making two cookies if no extra effort is involved on the agent's part. It's not just that the homeless people are only marginally happy, but that the agent is gratuitously preventing that happiness.

However, the implications are much worse than simply preventing goodness. Suppose that the absolute level of ALSC was at 810 and that the current state of affairs was ten times that. Under ALSC an agent is allowed to commit gratuitous murder so long as the outcome of the act does not dip below 810. For example, if the current level was

8,100 and murdering an innocent child was -2,000, that murder would be permissible since the outcome would be 6,100, well above the absolute level of 810. Attempts to get around this by, for example, moving to an individual situational absolute level satisficing consequentialism (ISALSC), as Bradley points out, have similar problems⁵. First, the gratuitous prevention of good objection still holds so long as everyone is above the absolute level. But if just one person is below the absolute level, then nothing is permissible until his level is elevated to at least the absolute level. This amounts to maximizing consequentialism. The world required for ISALSC to be intuitively plausible would be one where the utility hovers slightly above the absolute level. Since this is an unrealistic requirement it isn't a desirable alternative to ALSC.

III

Bradley points out one version of SC, namely self-sacrificing satisficing consequentialism (SSSC), which seems to get around these objections. Garret Cullity formulates it thus:

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.⁶

Bradley refers to this as Cullity's self sacrificing absolute level consequentialism (CSSALSC). Bradley says "CSSALSC does not allow for gratuitous prevention of goods. I suspect it...might be the most promising version of satisficing consequentialism"⁷.

However, Bradley does say that CSSALSC has some problems. He gives the following example to show one of the problems:

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:⁸

		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

He asks us to suppose that $n=20$. This makes A2 permissible since it meets the lower limit and both A1 and A3 are worse for Jack. Bradley says that this still prevents a better outcome and that seems implausible, but notice that it does not arbitrarily allow A2. A2 is allowed because it meets the threshold and it's better for Jack. Whereas other forms of SC gratuitously permitted the prevention of goodness, under CSSALSC it is not gratuitous at all.

Bradley's other objection is that it seems wrong to say that both A2 is permissible and A3 impermissible since doing A3 is better than A2. This does seem wrong, if the viewpoint is from non-agent-relative consequentialism, but the self-sacrificing aspect of CSSALSC is an agent relative view meant to give agent-centered options so this objection seems unclear. Is Bradley saying that there shouldn't be agent-centered options? If that's the worry then the objection should just be that the move to CSSALSC is unmotivated. That is a valid concern. In fact Norcross points out that "[t]he classical utilitarians all endorsed a non-agent-relative standard for assessing actions" and this might be a reason to reject agent-relative standards⁹. But CSSALSC is not meant to appeal to someone who rejects agent-relative standards. It is meant to appeal to someone

who thinks that agent-centered options are needed in consequentialism so as to bring it more in line with common sense morality.

If our motivation is to bring consequentialism more in line with common sense morality and more specifically we want to reduce the demandingness of consequentialism by giving agent-centered options, then CSSALSC seems to accomplish just that.

Furthermore, it does so without allowing the gratuitous prevention of goodness. It also gives us the right results when compared with common sense morality. Consider this example: A mortgage broker has found two mortgages that her client is eligible for. The first mortgage has a rate of 4.75% and the second a rate of 4.95%. The commission paid to the broker is the same for both loans. Suppose the lenders are so wealthy the benefit to them is minute. Lastly, let's assume the lower limit, n , is 20. The utility works out like this

	Interest Rate	$U_s(a)$	$U_{-s}(a)$	$U(a)$
A1	4.75%	50	100	150
A2	4.95%	50	-100	-50

Where $U_s(a)$ is the utility for the agent, $U_{-s}(a)$ is the utility for others and $U(a)$ is the overall utility. Both TAC and CSSALSC get this right. The broker should give the borrower the better interest rate. A2 is impermissible under TAC because there is a better alternative, and it is impermissible under CSSALSC since there is an alternative that is better overall and not worse for the broker.

Now consider this example. There is another option available at a rate of 4.85% that doubles the broker's commission. This new example works out as follows:

	Interest Rate	$U_s(a)$	$U_{-s}(a)$	$U(a)$
A3	4.75%	50	100	150
A4	4.95%	50	-100	-50
A5	4.85%	100	-50	50

Again CSSALSC gives us the right results. Now both A3 and A5 are permissible. A5 is permissible because it has at least a utility of 20 and the better alternative is worse for the broker. But on TAC A5 would be impermissible since there is an option with more overall utility.

Finally let's consider an example where the threshold isn't met. Let's suppose that the commission for A5 in the previous example wasn't double but rather only 32% more. Let's further suppose the utility worked out thus:

	Interest Rate	$U_s(a)$	$U_{-s}(a)$	$U(a)$	CSSALSC
A6	4.75%	50	100	150	Permissible
A7	4.95%	50	-100	-50	Impermissible
A8	4.85%	66	-50	16	Impermissible

In this example, even though we consider the agent's sacrifice for A6 and A7, we do not allow her to perform A8. A8 satisfies the first conjunct of the condition, that is that all the other alternatives are worse for the broker, but not the second conjunct, i.e. that A8 has a utility of at least 20. Again, on TAC only A6 is permissible since it is the maximizing alternative. These examples show how satisficing and self-sacrifice work together to provide intuitive results more in line with common sense morality than TAC.

One objection might be that satisficing isn't needed and that the theory would be more concise without it. So the theory might look something like Portmore's formulation of Schefflerian Utilitarianism (SU)¹⁰:

SU: S's performing [a] is morally permissible if and only if there is no available act alternative that would produce both (i) more utility for others (i.e., those other than S) than [a] would and (ii) at least as much egoistically-adjusted utility, where we include everyone's utility but adjust the overall total by giving S's utility ten times the weight of anyone else's.

Considering the following example with this formulation:

	Interest Rate	$U_s(a)$	$U_{-s}(a)$	$U(a)$	$U_{+s}(a)$	SU
A9	4.75%	50	100	100	600	Permissible
A10	4.95%	50	-100	-100	400	Impermissible
A11	4.85%	66	-50	16	610	Permissible

Where $U_{+s}(a)$ is $U_{-s}(a) + 10 \times U_s(a)$. A9 is permissible under SU because there is no other act that has more utility for others, therefore falsifying the first conjunct. A11 is permissible because there is no act with at least as much egoistically-adjusted utility, which falsifies the second conjunct. SU, like CSSALSC, provides agent-centered options, but with very different results. In this example this could be remedied by modifying the multiplier of ten in the formula, but we should examine more important differences between the theories.

Besides the multiplier for the agent's utility, the major difference between CSSALSC and SU is satisficing. We might ask what satisficing bought us and whether it was actually needed. In CSSALSC satisficing is meant to produce a maximizing consequentialist view when the circumstances are such that the absolute level cannot be reached. For instance, suppose that someone is about to have a tooth pulled and the dentist has the option to either give the patient nitrous oxide (N₂O) thereby relieving the

patient from suffering any pain during the procedure, shooting the patient up with Novocain and thereby reducing most of the pain experienced during the procedure, or doing nothing at all. Suppose N₂O costs the dentist five times the cost of Novocain, but doing nothing costs the dentist nothing. Further suppose the utility works out as such:

	Relief	$U_s(a)$	$U_{-s}(a)$	$U(a)$	$U_{+s}(a)$	SU	CSSALSC
A12	N ₂ O	-500	0	-500	-5000	Permissible	Impermissible
A13	Novocain	-100	-50	-150	-1050	Permissible	Permissible
A14	None	0	-500	-500	-500	Permissible	Impermissible

This demonstrates an important distinction between absolute satisficing and non-satisficing views. Since none of the acts meet CSSALSC’s lower limit of 20, the agent must perform the maximizing alternative. However, under SU the agent can perform any of the available alternatives. In this case it seems like SU gets the right outcome.

To further demonstrate why, imagine a fourth alternative (A15) that cost the dentist nothing but caused the patient 151 dolars. Is it really wrong for the dentist to select that option? If it is, then are we really providing agent-centered options? If this is a genuine worry, that the artificial threshold of 20 will restrict an agent’s options, then the following might appeal to us. Rather than using an absolute level of satisficing we might consider a relative level of satisficing. Let’s call it self sacrificing relative satisficing consequentialism (SSRSC) and it might be that:

SSRSC: there is a number, n ($n > 0$), such that: An act, a , performed by agent S, is morally right iff its utility plus n is greater than or equal to the utility of a maximizing alternative, and any better alternative is worse for S than a .

For the following examples assume $n = 20$. This allows the threshold to fluctuate relative to the maximizing alternative and should in turn produce more intuitive results. SSRSC produces the same results as CSSALSC in the table above, but for the fourth alternative I

mentioned, A15, the result is that it is permissible under SSRSC. It is permissible because the utility -151 plus the lower limit 20 equals -131 which is at least as high as the maximizing alternative (-150).

An important difference to note is that the satisficing versions set limits relative to the act and SU sets limits relative to the agent. Considering Norcross' worry regarding the statement Bentham made not to count any person for more than one, satisficing might be the more palatable view. While self sacrificing satisficing consequentialism (SSSC) views allow an agent to perform an act that doesn't maximize utility, it does not, as SU does, count the agent's value to be more than any other person. This is a delicate distinction as it seems like we *are* weighting the agent's value to have a greater weight in SSSC, but it's not the weight that's greater, rather the method of determining permissibility that allows agent-centered options. This distinction is clear when viewing a table for an SSSC view compared to SU. The SSSC view simply extracts the individual's utility from the overall, whereas SU multiplies that view by a factor of ten thereby counting that individual as much more than one.

V

The last question to consider is whether satisficing is necessary to provide a coherent consequentialist theory that has both self sacrificing options and self favoring options. I do not intend to definitively answer this question, but I hope to come away with some clarity regarding the move to satisficing. As I have shown, SU provides similar options as satisficing, but weighting the utility of the agent seems incongruent with, what I take to be, a traditional appeal of consequentialism, i.e. each person counting for one

and no more. It seems this last question might be cleared up by examining whether common sense morality follows the ranking for weighted self interest more or satisficing more. To find the answer we can turn to some ranking examples.

In our first example we will take the case of the inn keeper who offers the stranded motorist a free room. Her options are:

	$U_s(a)$	$U_{-s}(a)$	$U(a)$	$U_{+s}(a)$	CSSALSC	SU
Offer nicest room	-20	300	280	100	Permissible	Permissible
Offer average room	-20	100	80	-100	Impermissible	Impermissible
Don't offer a room	0	-100	-100	-100	Impermissible	Impermissible

In this case the results seem consistent with TAC, and I think someone inclined to a consequentialist theory would agree. How about the example of the child drowning in a shallow pond? It might work out something like this:

	$U_s(a)$	$U_{-s}(a)$	$U(a)$	$U_{+s}(a)$	CSSALSC	SU
Save the child	-50	1000	950	500	Permissible	Permissible
Call for help	-20	800	780	600	Permissible	Permissible
Keep walking	0	-1000	-1000	-1000	Impermissible	Impermissible

This example differs from TAC in that it gives options, but we notice that CSSALSC and SU both provide the correct result. Now we will take an atypical example that will show where CSSALSC and SU come apart:

	$U_s(a)$	$U_{-s}(a)$	$U(a)$	$U_{+s}(a)$	CSSALSC	SU
A16	100	-350	-250	650	Impermissible	Permissible
A17	-20	800	780	600	Permissible	Permissible
A18	50	50	100	550	Permissible	Impermissible

It does not seem to be in line with common sense morality that an agent can cause so much suffering to others for his own benefit as SU allows in A16. Imagine the agent is a

sadist, should it be permissible for him to cause others suffering for his own benefit? CSSALSC isn't immune to this objection, but it does provide a stronger formula for avoiding large disparities between self and others. And as mentioned in the previous section, a formulation such as SSRSC might provide even more intuitive result than the absolute level theory does.

Conclusion

This paper was meant to show the motivation toward a self/other consequentialist theory. There are varying approaches each with their own obstacles. I have contrasted two views here with a special emphasis on satisficing. Although Bradley, Portmore, Pettit and others have argued against satisficing consequentialism, it seems like this theory deserves another look. Especially as we try to move consequentialism more in line with common sense morality.

I have shown that satisficing is motivated by the need for self-sacrificing and self-favoring agent-centered options. Furthermore, I have shown how it holds up compared to other self/other theories. Satisficing has its own limitations, but I have shown that it is not so much more than some of the other options available to satisfy the move to agent-centered options. This paper was not meant to show that satisficing was a superior option, rather the intention was to leave satisficing available as an option to satisfying the move toward a theory more in line with common sense morality.

Notes

1. Douglas Portmore makes the case that Schefflerian Utilitarianism is a better option in *Notes for "Consequentialism and the Self-Other Asymmetry"*
2. Michael Slote, *Satisficing Consequentialism*
3. Michael Slote, *Satisficing Consequentialism* (p 147)
3. Michael Slote, *Satisficing Consequentialism* (p 142)
5. Ben Bradley, *Against Satisficing Consequentialism* (p 100)
6. Ben Bradley, *Against Satisficing Consequentialism* (p 107)
7. Ben Bradley, *Against Satisficing Consequentialism* (p 107)
8. Ben Bradley, *Against Satisficing Consequentialism* (p 107)
9. Alastair Norcross, *Good and Bad Actions* (p 12)
10. Douglas Portmore, *Notes for "Consequentialism and the Self-Other Asymmetry"* (p 3)