



Slater, J. (2019) Satisficing consequentialism still doesn't satisfy. *Utilitas*,
(doi:10.1017/S0953820819000402).

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/201179/>

Deposited on: 21 October 2019

Enlighten – Research publications by members of the University of Glasgow_
<http://eprints.gla.ac.uk>

Satisficing Consequentialism Still Doesn't Satisfy

JOE SLATER

University of Glasgow

Abstract

Satisficing consequentialism is an unpopular theory. Because it permits gratuitous suboptimal behaviour, it strikes many as wildly implausible. It has been widely rejected as a tenable moral theory for more than twenty years. In this paper, I rehearse the arguments behind this unpopularity, before examining an attempt to redeem satisficing. Richard Yetter Chappell has recently defended a form of “effort satisficing consequentialism”. By incorporating an “effort ceiling” — a limit on the amount of willpower a situation requires — and requiring that agents produce at least as much good as they could given how much effort they are exerting, Chappell avoids the obvious objections. However, I demonstrate that the revised theory is susceptible to a different objection, that the resulting view requires that any supererogatory behaviour must be efficient, which fails to match typical moral verdicts.

In this paper, I claim that understanding what is morally ‘good enough’ in terms of effort rather than the goodness of outcome produced cannot salvage satisficing consequentialism. Traditional versions of the theory have suffered extensive criticism, largely due to their condoning of gratuitous preventions of goodness. In the first section, I present the motivations for

moving towards satisficing consequentialism, and rehearse the arguments against that move.

The second section acknowledges an attempt to change this. Richard Yetter Chappell has recently attempted to resurrect satisficing consequentialism.¹ By incorporating an “effort ceiling” – a limit on the amount of willpower a situation requires – but requiring that agents produce *at least* as much good as they could given how much effort they are exerting, Chappell avoids the common objection.

Finally, I argue that Chappell’s account leaves his satisficing consequentialism susceptible to another criticism. His position holds that any supererogation *must* be perfectly efficient. While there are notable cases where doing nothing is permissible, but doing some (but not *the most*) good is not, many examples do not seem to have this feature. I provide several such examples. I suggest that this is a very serious problem for an effort-based view of satisficing consequentialism, and indicates that any such view is still unsatisfactory.

1 The Rise and Fall of Satisficing

Not content with the existing forms of consequentialism, Michael Slote popularised (but later rejected) satisficing consequentialism.² Typical act-consequentialism, which requires that one bring about the *most* good, seems extremely demanding. If correct, many of us are under a moral obligation to work extremely long hours, perhaps in jobs that do not fulfil us (if unfulfilling jobs pay significantly more), and donate the vast majority of our disposable income to the best causes. We could spend enough on ourselves

¹Richard Yetter Chappell, ‘Willpower Satisficing’, *Noûs* 53 (2019), pp. 251-263.

²Michael Slote, ‘Satisficing Consequentialism’, *Proceedings of the Aristotelian Society*, suppl. vol.58 (1984), pp. 139-163.

to ensure that we are psychologically able to continue this sort of life, but our lives would be very austere.

Slote recognised that these conclusions about how demanding morality is seem to defy common-sense.³ Noting the appeal of viewing rightness as depending on the consequences of an action, he thus sought to soften the theory. No longer would the *most* good be required. Instead, one could permissibly just do “enough”. Abandoning the maximising constraint does seem like a sensible move for those sympathetic to the consequentialist project.

The moral satisficer claims that agents may, sometimes, elect to perform an action that is good enough, regardless of whether it is the best moral option. The claim is not that that a decision procedure that does not require maximising is actually best in the long run. That may be a plausible claim – efforts to do the best might be stressful, or lead to an agent failing to achieve a good outcome due to refusing to settle for less than perfect, and ending up with nothing – but that is not the focus here.⁴ Satisficing consequentialists accept what Tim Mulgan calls “blatant moral satisficing”, the claim that even “when the agent already knows which particular act is best from a Consequentialist perspective, she is still perfectly justified in selecting any other act which is good enough”.⁵ It holds that it is permissible to act sub-optimally in cases where there is no moral reason that speaks in favour of doing the less good option.

Slote illustrates this with two main examples. First, he considers a hotel manager who discovers a family whose car has broken down outside the

³Slote, ‘Satisficing Consequentialism’, p. 142.

⁴Even non-consequentialists might endorse some form ‘strategic satisficing’. Christine Swanton, for instance, suggests that strategies to satisfice rather than pursue the good directly do seem rational, e.g. Christine Swanton, ‘Satisficing and Virtue’, *The Journal of Philosophy* 90,(1993), pp. 33-48.

⁵Tim Mulgan, *The Demands of Consequentialism*, (Oxford, 2001), p. 128.

building.⁶ They do not have money to pay for a room or a meal, but the manager offers them one of the rooms in the hotel out of beneficence. She does not offer them the best room, but simply the first suitable one she finds available, which she deems to be good enough, knowing that there are other better rooms she could find very easily.

Second, he provides an example of a fallen warrior. Because of his heroism, the gods decide to grant him a single wish for those he is leaving behind.

Presented with such an opportunity, may not the warrior wish for *his family* to be comfortably well off forever after? And will we from a common-sense standpoint consider him to have acted wrongly or non-benevolently towards his family because he (presumably knowingly) rejected an expectably better lot for them in favour of what was simply good enough? Surely not.⁷

Slote suggests that we would not condemn the hero not wishing that they be only “comfortably well off” rather than extremely well off. Even though alternative wishes would have required the same effort, Slote claims that the less generous wish is permissible.

This allows the satisficer to easily evade the demandingness objection. Agents may permissibly give some of their income to successful charities, and keep a nice amount for themselves, so long as their actions are good enough.

There are a variety of ways satisficing consequentialism can be understood. These typically vary with how the “good enough” condition is interpreted. Slote actually suggests that the “enoughness” will be formalisable in terms of some “percentage or mathematical function” taking into account

⁶Slote, ‘Satisficing’, p. 149

⁷Slote, ‘Satisficing’, pp. 150-151.

of the amount of good that one could do.⁸ This would permit agents to act in many ways we do deem praiseworthy, yet still forbid complete moral slacking off. When an agent *does* do more than is “good enough”, they act supererogatorily. There are many different ways that “good enough” could be interpreted. Ben Bradley, for instance, offers *six* different ways.⁹ All of the views, however, fall prey to the same devastating objection.¹⁰

The objection is due to the permissive verdicts satisficers must give about gratuitous prevention of goodness. The typical moral satisficer must accept that it is sometimes permissible to perform an action worse than one could. In order to respond to the demandingness objection, a huge motivation (if not the entire *raison d'être*) for satisficing views, it must actually permit actions *much* worse than maximising. To see this, we can think about the type of life that the satisficer might want to think is permissible. Even if an agent gives considerably more than average to an effective charity (say 40% of a modest income), and spends three nights a week volunteering, as well as being a generally friendly, kind person in their normal interactions, they are falling *way* short of maximising. But surely this is precisely the type of behaviour a satisficer wants to say *is* permissible. So, satisficers must accept that acting significantly sub-optimally is in some (many, in fact) situations permissible.

But once we accept this, very counterintuitive verdicts follow. Tim Mulgan illustrates this with the following example:

The Magic Game: Achilles is locked in a room, with a

⁸Slote, ‘Satisficing’, p. 156.

⁹Ben Bradley, ‘Against Satisficing Consequentialism’, *Utilitas* 18 (2006), pp. 97-108, at 101.

¹⁰Bradley describes how each of the six interpretations are understood, and how each of them cannot respond adequately to the objection (‘Against Satisficing’, pp. 101-104). I will not repeat these explanations here.

single door. In front of him is a computer screen, with a number on it (call it n), and a numerical keypad. Achilles knows that n is the number of people who are living below the poverty line. He also knows that, as soon as he enters a number into the computer, that many people will be raised above the poverty line (at no cost to Achilles) and the door will open. There is no other way of opening the door. Because of the mechanics of the machine, any door-opening number takes as much time and effort to enter (negligible) as any other.¹¹

Achilles then inputs some number (p) lower than n . Obviously, in doing so, he condemns $n - p$ people to continued poverty. His only justification for doing so is that this number was “good enough”. Our response to this sort of gratuitous prevention of good is obviously unfavourable, and for most of us, this is a knock-out blow for (traditional) satisficing consequentialism.¹²

Satisficing consequentialism is supposed to help us avoid some of the (seemingly) bad verdicts we get from maximising act consequentialism, namely that we act wrongly almost always, and that what seem like generous people actually (except when they do successfully maximise) act impermissibly. But if we accept satisficing, it looks like we must condone behaviours like Achilles’s here. However, this result seems (for many of us) similarly unpalatable.¹³

Because of this objection, the move to a satisficing view seems fruitless.

¹¹Tim Mulgan, ‘Slote’s Satisficing Consequentialism’, *Ratio* 6, pp. 121-134, at 125.

¹²Elsewhere, Mulgan offers even more alarming examples, suggesting that satisficers are committed, in some circumstances, to condoning murder (Tim Mulgan, ‘How Satisficers Get Away With Murder’, *International Journal of Philosophical Studies* 9 (2001), pp. 41-46).

¹³de Lazari-Radek and Singer note that this conclusion “may be just as hard to accept as the maximizing conclusion” delivered by act consequentialism (e.g. Katarzyna de Lazari-Radek and Peter Singer, *The Point of View of the Universe: Sidgwick and Contemporary Ethics* (Oxford, 2014), p. 143.)

As Bradley correctly concluded, to be made acceptable, satisficers must demonstrate both that their theory is well-motivated (i.e. that there is some problem it solves that extant views fail to deal with appropriately) *and* that they are able to provide a version that does not permit the gratuitous prevention of goodness. In the next section, I discuss the new satisficing account Chappell has formulated that appears promising.

2 Satisficing Consequentialism: The Reboot

Act consequentialism is often rejected because it seems to demand too much. If the demandingness is understood simply in terms of costs to welfare, the objection fails because act consequentialism actually minimises costs to welfare, once the welfare of the moral patients is taken into account.¹⁴

Due to this type of worry, Chappell and others have speculated that we should instead understand what makes an action too demanding in terms of *difficulty*.¹⁵ Chappell in particular discusses the willpower an action requires. For this reason, he describes his view as *Willpower Satisficing Consequentialism*.¹⁶

Chappell holds that agents are morally required to act in a way that displays adequate moral concern. In explicating his view of moral concern, Chappell draws upon “Quality of Will” accounts. These hold “feelings of

¹⁴This problem for demandingness objections is discussed by David Sobel (‘The Impotence of the Demandingness Objection’, *Philosophers’ Imprint* 7, pp. 1-17), and leads him to reject that the demandingness objection against act consequentialism has any force.

¹⁵For further discussion of this conception, see McElwee (Brian McElwee, ‘What is Demandingness?’, in M. van Ackeren and M. Kuhler (eds.), *The Limits of Moral Obligation: Moral Demandingness and Ought Implies Can* (New York, 2016), pp. 19-35), who is favourable, and Cohen, who briefly discusses but dismisses it (G. A. Cohen, ‘If You’re an Egalitarian, How Come You’re So Rich?’, *The Journal of Ethics* 4 (2000), pp. 1-26, at 24.). Dana Nelkin also defends the view that moral status of an action (the degree of praiseworthiness/blameworthiness) is dependent upon its difficulty (Dana Nelkin, ‘Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness’, *Noûs* 50 (2016), pp. 356-378.

¹⁶Henceforth WSC.

blame to be warranted when the target of the attitude has acted in a way that displays inadequate moral concern.”¹⁷ Specifically, Chappell understands “moral concern” as a combination of two factors:

Altruistic Concern: The extent to which an agent values others’ welfare.

Strength of will: The mental effort an agent puts in to promoting general welfare.

Someone with limited altruistic concern – a miserly, Scrooge-like character – thus may have to expend a great deal of mental effort to express the adequate moral concern a much more agreeable person would express easily.

Rather than outcomes, it is the moral concern an agent expresses which must be good enough. For Chappell, the level of moral concern which is good enough is settled by our theories of blameworthiness, and specifically, quality of will accounts. Greater moral concern typically promotes the general good and it is obligatory to express *sufficient* moral concern. More is better, but less is impermissible.

Because what moral concern counts as sufficient is settled by our theories of judgement (e.g. about the quality of will), and the moral concern a person exhibits is a function of their altruistic concern and their strength of will, Chappell is able to express the maximum effort an agent must expend in a situation. He calls this the *effort ceiling*, which he defines as follows:

The Effort Ceiling X (for an agent at a time): that value X such that, *in order to qualify as “adequately concerned”* according to our prior the-

¹⁷Chappell, ‘Willpower Satisficing’, p. 256. Specifically Chappell makes use of the accounts given by Arpaly (Nomy Arpaly, *Unprincipled Virtue: An Inquiry Into Moral Agency* (Oxford, 2003)) and Strawson (Peter Strawson, ‘Freedom and Resentment’, *Proceedings of the British Academy* 48 (1962), pp. 1-15). I will not discuss the intricacies of these accounts here.

ory of blameworthiness and quality of will, the agent must be willing to expend up to X willpower in pursuit of the general good.¹⁸

To elucidate, we have some prior notions about how much people need to do with regards to promoting the general good. Moral concern is seen as a combination of both altruistic concern *and* strength of will. A particularly selfish agent will have minimal altruistic concern, so to meet X , they may need to put in a great deal of willpower, whereas someone who is particularly other-regarding (having a high amount of altruistic concern) may need a very low amount of willpower to meet X . It does not take much willpower on behalf of an altruistically-disposed agent in order for them to act well, whereas the miser may find that it extremely motivationally difficult.

Chappell makes use of this effort ceiling in order to offer his version of willpower satisficing consequentialism.

WSC An act A is permissible iff it produces no less utility than any alternative action the agent could perform at the cost of up to either (i) X willpower, or (ii) the willpower cost of A , whichever is greater.¹⁹

Any act which causes less utility than could be produced with X , or which does not produce as much utility as could be provided for the amount of effort expended, is thus regarded as impermissible. The limits of moral demands are set by the willpower effort ceiling. With this in mind, we can see that a prescription is overly demanding if it requires more willpower than X .

As Chappell notes, this gives felicitous verdicts in the cases of gratuitous prevention of goodness. The standard moral satisficer must accept that it is

¹⁸Chappell, 'Willpower Satisficing', p. 257.

¹⁹Chappell, 'Willpower Satisficing', p. 255.

sometimes permissible to perform an action *much* worse than one could, at a whim or for no reason, and even though it is just as easy – requires just as much exertion of the will – as other options. Chappell is able to avoid this unpleasant consequence. Consider Mulgan’s Magic Game example. The effort ceiling, X , in this case, is the (minimal) effort of Achilles inputting the numbers. WSC then requires that he do as much good as possible, with this amount of effort (or perform some morally efficient act which does more good). As we would hope, WSC will give us the verdict that Achilles must input the number n .

Mulgan actually suggested that an “effort satisficing” view – like Chappell’s – was possible in *The Demands of Consequentialism*, noting a temptation to “think that Slote can construct a Satisficing Consequentialism based on the notion of ‘good-enough effort’”.²⁰ However, this is definitely not what Slote had in mind. In fact, Chappell’s account could not give the verdicts that Slote favours in several of his examples. This might not be too much of a problem for WSC, as these cases (the Hotel case and Fallen Hero) are not the sort that motivate satisficing consequentialism (i.e. instances where the demands placed on the agent by act consequentialism seem unreasonable).

Mulgan raises a potentially bigger worry, that Slote could not make the ‘good enough’ concern effort rather than outcomes, because it would not provide a *consequentialist* explanation for the permissibility of sub-optimal acts:

[Slote’s] original theory provides an explanation: morality does not make unreasonable demands, because it never requires that agents produce more than a certain amount of good. Here, the limit on the costs an agent can be asked to bear is explained

²⁰Mulgan, *Demands of Consequentialism*, p. 138.

by appealing to a feature of the outcome of her actions. This is both an *explanation* and a *Consequentialist* explanation.²¹

However, if the reason we are permitted to act sub-optimally is because it is accepted that we are not required to expend unreasonable efforts, this is no longer justified, at its core, in terms of consequences. A major motivation for satisficing consequentialism is the general appeal of consequentialism but a dislike of the maximising aspects, so if the resulting theory is no longer consequentialist – explaining obligations in terms of values of outcomes – this might be seen as a significant cost for WSC.

I don't think this is a major concern. The WSC formulation *does* define rightness in terms of utility, so it does seem appropriate to call this consequentialist.²² My worry, however, is not that the theory isn't consequentialist enough, but that it still fails to give the right verdicts. In the next section I suggest that Chappell's account incorrectly assesses as impermissible actions which are good, but not *as good* as others which require the same effort.

3 Responding to Chappell

A feature of Chappell's account is that it requires that any supererogation be efficient. While this is accurate in some cases, I argue that it overgeneralises. Sometimes, an agent can do more than required without having to do as much good as possible for the willpower exerted.

²¹Mulgan, *Demands of Consequentialism*, p. 138.

²²In one sense I am not interested in whether it should be categorised in this way – if it's right, it's right! But in terms of what *explains* the moral properties of acts – as suggested in Gary Watson, 'On the Primacy of Character', in O. Flanagan and A. Rorty (eds.), *Identity, Character and Morality: Essays in Moral Psychology*, (Cambridge Massachusetts, 1997), pp. 449-469 – this does seem appropriate.

Of course, we can grant that acting in a way to deliberately and gratuitously cause less good than one might is often a bad thing. This is clear in Mulgan's Achilles case. We can also notice this in cases where an agent seems to be under no obligation to act at all. Consider Kagan's case where an agent has to choose whether or not to run into a burning building to rescue survivors.²³ We may accept, given the level of risk that they would face, that it is perfectly permissible for an agent not to enter, i.e. to do nothing (expending zero effort). However, if they do choose to enter, and once inside they have a choice to save either a bird or a baby, it seems that they are under an obligation to save the baby, rather than the bird. Doing nothing is fine, running in to save the baby is fine, but running in to save the bird (we typically think) is not okay. This is an instance of gratuitous sub-optimal supererogation.

Chappell is also able to get the correct verdict here. If one is under no obligation, the X value would be zero, but if one expends the effort to run into the building, they must produce no less utility *than any alternative* they might perform at the cost of up to that willpower. Again, WSC looks good. However, while WSC successfully condemns gratuitous prevention of goodness, it also condemns any supererogation which is inefficient, i.e. that does not produce the most possible utility for the effort expended. This causes some problems, as not all instances of inefficient supererogation seem morally pernicious.

Consider philanthropic giving. While squandering one's money by donating to a charity one believes foolishly wastes their donations (but does *some* good) strikes us as clearly bad, many of us would recoil at the prospect of condemning anyone who gives a substantial amount of their income to a

²³Shelly Kagan, *The Limits of Morality* (Oxford, 1989), p. 16.

not-perfectly-efficient charity – perhaps even the second-best charity. Chappell’s account requires that any time we go above the effort ceiling, we are morally required to do the best we could at that effort level. This seems implausible. If someone did give all their disposable income to the second best charity, it would seem extremely inappropriate to condemn them for doing so (discarding any other bad-making moral features).²⁴

We can even consider much more mundane examples. Imagine the minimum you can do for Mothers’ Day is send a card and make a phone call. You realise also that the *best* thing to do, would be to take her out to her favourite restaurant (Mother’s Choice). However, a close-second would be to take her out to a restaurant you prefer (Child’s Choice), which she also likes. The amount you prefer the Child’s Choice (by stipulation) is slightly less than the amount your mother prefers Mother’s Choice. Assuming attending either restaurant would take the same amount of effort – the price of the food is similar, the distance to travel is the same, etc. – WSC is committed to condemning taking her to Child’s Choice. But you are permitted to simply call and send a card, which generates much less utility!

In this case, there is *something* that counts in favour of performing the act with less utility. *You* prefer the food at this restaurant. We might think that your knowing this could make it the case that going to Mother’s Choice is *slightly* more difficult – motivating oneself to go there takes a little bit more of your willpower. Unfortunately, if an agent’s mild preferences are taken to affect willpower in this way, WSC loses its ability to give the felicitous verdicts in many of the previous problem cases for satisficing con-

²⁴Even many effective altruists accept this verdict. Pummer, for instance, defends the claim that it is wrong to do *much* less good than one can when costs are equal, but accepts that in various circumstances it is permissible to do *some* less good (e.g. Theron Pummer, ‘All or Nothing, but If Not All, Next Best or Nothing’, *Journal of Philosophy* 116 (2019), pp. 278-291).

sequentialism. Take the bird and the baby case. If an agent has a particular fondness of birds which would make it *very slightly* more unpleasant for them to leave the bird, they would be permitted, according to WSC, to opt to save the bird over the baby. This is a very undesirable verdict.

Chappell could respond to examples like these by amending the second clause of WSC, which says that for actions requiring effort more than the effort ceiling, the act must promote as much utility as possible at that willpower cost. He might, for instance, remove it:

WSC² An act A is permissible iff it produces no less utility than any alternative action the agent could perform at the cost of up to X willpower.

However, doing so would lead a big problem, as this would *condone* all sub-optimal supererogation, even where this looks dubious, as with Kagan's bird-saver. Another example where this looks like the wrong result is in Joe Horton's illustration of the 'all or nothing problem', where an agent has the following options:²⁵

1. Do nothing
2. Save one child by allowing both your arms to be crushed
3. Save two children by allowing both your arms to be crushed

Because of the huge cost involved, it seems permissible to do nothing here, so for WSC no effort is required. If WSC was amended to only require actions that produce the maximum utility at that level of effort (none), any of the other two actions would be permissible. However, Horton contends, plausibly, that it is obviously wrong to decide to only save one child, when one could – at no extra cost or effort – have saved two.

²⁵Joe Horton, 'The All or Nothing Problem', *The Journal of Philosophy* 114, pp. 94-104, at 94.

Another way Chappell might amend WSC would be as follows:

WSC³ An act A is permissible iff (i) it produces no less utility than any alternative action the agent could perform at the cost of up to X willpower or (ii) A requires more willpower than X *and* satisfies condition C.

Where C stipulates a certain relationship between, for instance, the additional utility gained *for an agent* and the utility gained *external to the agent*. This could function like Scheffler's agent-centred prerogatives,²⁶ so that above the effort ceiling an agent could count their own interests as more weighty than others' interests. This could even be supplemented with constraints on harming others to avoid verdicts such as it being permissible to expend effort above the effort ceiling to harm someone else by a similar amount it would benefit an agent by.

However, whatever way such a condition, C, is construed, if it only operates over values of personal and impersonal utility produced, it still seems unsatisfactory. To see why, consider the following two cases.

1. Take Kagan's example of the bird and the baby, however in this case the bird is *mine*. Because of the danger it is permissible to do nothing. It is also permissible for me to run in to save the baby. It is not permissible to save the bird. Even if the bird is *mine*, saving the bird seems impermissible. Thus, the condition should not deem it permissible to make myself *slightly* happier to expend some effort in saving a bird rather than a baby.
2. I donate a significant sum of money (supererogatorily) to the second-best charity. I opt for this because I know someone who suffers from a

²⁶Samuel Scheffler, *The Rejection of Consequentialism*, (Oxford, 1982).

disease it treats, and this gives me a warm feeling inside. We typically view this as permissible.²⁷

Let me stipulate that the warm feeling in case 2 is, in terms of my personal utility, the same as I would gain from saving the bird in the first case. The difference in the alternative costs, however, is significantly higher in the charity example. The second-best charity might be able to save hundreds of lives fewer with the money I donate. So any condition simply operating over utility values could not allow WSC to give the desired verdicts.

Perhaps a more promising way of amending WSC could be to replace the second clause of WSC with something like an ‘Avoid Gratuitous Worseness’ principle.²⁸ This could impose a requirement that any inefficient giving is not simply *gratuitous*, i.e. that something counts in its favour. If this was not sufficiently strong, however, this would permit one to save the bird rather than the baby, or even save only one child’s life in Horton’s case, so long as the agent can find *some* reason that counts in its favour. Developing an account of what makes certain acts objectionably *gratuitous* would then be challenge. However, a satisfactory answer to this question looks set to render the rest of the WSC apparatus redundant, as this would seemingly tell us everything that WSC sought to explain.

Another response available is accepting that WSC clashes with some common intuitions, but still find it suitable. We could suggest that intuitions are misleading in certain situations, so it is reasonable to expect some verdicts that clash with our judgements. In the philanthropy case, effective

²⁷Maximisers like de Lazari-Radek and Singer would disagree (e.g. Katarzyna de Lazari-Radek and Peter Singer, ‘Doing our Best for Hedonistic Utilitarianism’, *Ethics & Politics* 18 (2016), pp. 187-207, at 197) but they reject the move to satisficing consequentialism anyway.

²⁸Theron Pummer, ‘Whether and Where to Give’, *Philosophy & Public Affairs* 44 (2016), pp. 77-95, at 93.

altruists and others might welcome the verdict that giving sub-optimally *actually is* morally impermissible.

However, I think we have independent reason to reject WSC, namely that it can condone certain actions justified self-interested reasons, but is unable to permit *better* actions, which would carry similar justifications. To illustrate, consider again the Mothers' Day case. Presumably, it is because I am allowed to give some weight to my own self-interested reasons that I am morally permitted to merely send a card and make a phone call. So self-interest permits me exerting less willpower, and consequently producing less good. WSC can accept this, as our blaming practices permit some self-interested behaviours. But, the same explanation – my self-interest – does not permit doing less good when I *do* want to put more effort in. Self-interest, under WSC, can therefore permit doing no good at all (or negligible good, in this situation), but not doing significant good inefficiently.

For WSC then, self-interest *can* justify omitting to act, but not justify choosing one option over another at the same willpower. It would seem bizarre if the permissibility of my acting slightly self-interestedly depends only on how much effort I exert.

4 Concluding Remarks

Chappell's willpower satisficing successfully avoids condoning the gratuitous prevention of good. This was a disastrous and unacceptable result. The examples I have provided show that WSC also gives strange verdicts. Denying the intuitive verdicts in the cases I have supplied, as required by WSC, would not require *as* revisionary treatment of our moral landscape.

But I do regard the verdicts as a problem, and a serious one. Chappell has, by implanting the 'quality of will' and willpower concepts, resurrected

satisficing consequentialism into a stronger candidate for a moral theory. However, the account still fails to match up with huge classes of typical moral intuitions.²⁹

²⁹I am grateful to two anonymous reviewers and the editors at *Utilitas* for helpful suggestions on a previous version of this paper. A short version of this paper was presented at the 2019 Joint Session at Durham, so thanks also to the organisers for the opportunity to present at that venue, and for the questions of the audience. Big thanks also to Lizzy Ventham for reading a typo-ridden and rambling version of this, and suggesting many much-needed cuts!