

SATISFICERS *STILL* GET AWAY WITH MURDER!

JOE SLATER

University of Glasgow

Recently, a few attempts have been made to rehabilitate satisficing consequentialism. One strategy, initially shunned by Tim Mulgan, is to suggest that agents must produce an outcome at least as good as they could at a particular level of *effort*. The *effort-satisficer* is able to avoid some of the problem cases usually deemed fatal to the view. Richard Yetter Chappell has proposed a version of effort-satisficing that not only avoids those problem cases, but has some independent plausibility. In this paper, I argue that we should be concerned by verdicts the effort-satisficer delivers that are too permissive. Revising a problem for the traditional outcome-satisficer, I argue that Chappell's willpower satisficing, and more generally, any effort-satisficer, must implausibly condone murder in many cases. This seems like a serious issue for any attempts at rehabilitating satisficing consequentialism in this way. After presenting my objection, I consider some ways an effort-satisficer might attempt to revise the account to avoid these problems, but argue that none of these can succeed.

SATISFICING consequentialism is a version of consequentialism designed to conform better with our intuitions than maximising versions. It promises to avoid overly demanding verdicts of our obligations, while maintaining many of the features found attractive about consequentialist views. However, satisficing views have been criticised as vulnerable to a series of problem cases, where they condone horrific acts, like gratuitous murders (Mulgan 2001a).

Most discussion of satisficing forms of consequentialism has been limited to what we might call *outcome-satisficing*, wherein agents are permitted to act in any way leading to an outcome that is deemed 'good enough', with various attempts to specify how this should be determined.¹ Despite a plethora of formulations

1. E.g., in Slote's original discussion, he mentions two ways this could be specified which he does not endorse (attributed to Bentham and Popper). Jamieson and Elliot's 'progressive consequentialism' (though they offer their view as an alternative to satisficing) offers a way of determining what is good enough based on whether agents make the world a better place (2009).

Contact: Joe Slater <joe.slater@glasgow.ac.uk>

being offered, one particular way of understanding satisficing consequentialism, *effort-satisficing*, has largely been neglected. On this view, what is required of agents is that their actions produce an outcome at least as good as they could at a particular level of effort, which must be determined by contextual features.²

Richard Yetter Chappell (2019) has recently defended a version of effort-satisficing consequentialism, Willpower Satisficing Consequentialism (WSC). Chappell argues that his version can avoid the cases that plagued outcome-satisficing consequentialism. If he is correct, WSC is a strong candidate for being the most plausible version of consequentialism.

However, I argue that WSC, and effort-satisficing views more broadly, are susceptible to a set of problem cases similar to those that affected the original view. I describe the structure of these cases and argue that they pose a significant problem. I consider two ways of amending the view, while maintaining the spirit of the original proposal, but argue that both either yield similar counter-intuitive verdicts or are no longer able to adequately respond to demandingness objections. Thus, as things stand, effort-satisficing consequentialism should be rejected.

I begin by depicting the differences between outcome-satisficing and effort-satisficing, and show how Chappell's recent manoeuvre is able to avoid condoning murder in the standard problem cases. I then describe the structure for my new problem cases, demonstrating that effort-satisficing gives implausibly permissive verdicts. Finally, I discuss two amendments, and illustrate why these are inadequate.

1. Effort-Satisficing

Act-consequentialists standardly hold that agents are morally required to act in a way that brings about the *best* consequences, but this has extremely demanding implications that many find implausible. According to this position, actions we typically view as morally benign, like going to the cinema, buying a spare pair of shoes or having a drink at the pub are actually morally impermissible, because our time or resources could do more good if used differently. In response to counterintuitive verdicts like this, Michael Slote proposed a *satisficing* form of consequentialism (1984). According to this view "all actions with good enough consequences are right" (Jenkins & Nolan 2010: 451), and sometimes it is the case

2. In this paper, I contrast effort-satisficing consequentialism with outcome-satisficing consequentialism. However, effort-satisficing should be understood as a species of outcome-satisficing, because it is still the value of the outcome which must be "satisficed", i.e., that needs to be good enough. The distinguishing feature of the effort-satisficer is that considerations of effort fix what counts as "good enough". I thank an anonymous reviewer for pressing me to clarify this taxonomical point.

that non-optimal consequences are good enough. Slote thought that this view would be superior to maximising views in terms of its “common-sense moral plausibility” (1984: 142).

Despite Slote’s hopes, satisficing consequentialism has been widely rejected since Tim Mulgan demonstrated its serious problem in permitting acts we ordinarily see as morally reprehensible (1993; 2001a). For traditional versions of the view, problems arise because however we understand what makes consequences ‘good enough’, they either seem to require us to be too restrictive, defeating the purpose of the amended version of consequentialism, *or* they permit too much, with implausible repercussions. To illustrate this dilemma, consider that we devise some function, evaluating the goodness of the consequences, such that if the act passes a certain test in the hedonic calculus (or the calculus of how-ever we are measuring the goodness of consequences), our account regards it as permissible. The function the satisficer provides may either permit agents to act sub-optimally in large classes of cases, or it may not. If it does not, it will not give us a significantly less demanding moral theory. As giving a suitable response to demandingness objections was the *raison d’être* for satisficing consequentialism, this is a significant failing. This is the first horn of the dilemma.

So they must make the function such that it permits agents to act sub-optimally in many everyday cases. However, if the function works like this, perhaps generating only 90% of the maximum good possible,³ it will permit acts that seem independently abhorrent.

For instance, Mulgan provides an example of a trolley case, where the best option one could take would be pushing a heavy sandbag to stop a trolley heading towards a cliff (2001a: 42). This would stop the trolley and save the ten people on board. However, one could also shoot a nearby person, Bob, causing him to fall onto the tracks. This would kill Bob, but save the ten people on board the trolley. This seems like it would generate 90% of the possible good (assuming each person would contribute similar amounts to the general good, and there are no other negative side-effects from the murder).

In many cases like this, an agent deliberately aiming for a less good option, particularly when significantly morally better outcomes could easily be achieved, seems morally outrageous. This problem, described by Bradley as “gratuitous prevention of goodness” (2006) is a huge problem for traditional forms of satisficing consequentialism. Bradley describes a host of different ways we could interpret the view,⁴ and describes ways each of these has the problem of permitting gratuitous prevention of goodness.

3. Slote did suggest that what was “good enough” could be determined by some “sort of percentage or other mathematical function” (1984: 156) of the overall good that could be produced.

4. Bradley considers many ways the outcomes could be relevant, e.g., how close they are proportionately to the optimal outcome, how good they are for each individual, and taking the

However, each of the varieties of satisficing consequentialism Bradley discusses can be described as *outcome-satisficing*. We can distinguish a distinct type of satisficing from the type that Bradley was concerned with. To illustrate compare the following:

Outcome (traditional) Satisficing: An act φ is permissible iff either φ maximises the good *or* the outcome of φ is *good enough*.

Simple Effort-Satisficing: An act φ is permissible iff either:

1. φ maximises the good
OR
2. φ generates *enough* good, where this is set by the amount of good the agent could bring about at effort level X.

The effort-satisficer must then give us some way of determining what amount of effort might be required in a certain situation (X).⁵

Mulgan acknowledged in *The Demands of Consequentialism* the possibility of adopting something like effort-satisficing, of a view where “the ‘good enough’ refers, not to the agent’s contribution, nor to the outcome resulting from that contribution, but to the effort required to produce that outcome” (2001b: 137) but dismissed it. He did not think the effort or difficulty involved in acting could provide an explanation (and particularly not a *consequentialist* reason) for the permissibility of suboptimal acts.

Recently, however, considerably more attention has been paid to effort and its moral relevance. Chappell argues that “it is the effort required rather than the prudential cost of the effort, that is the proper measure of the burden placed

outcome of the agent into special consideration, but it is always only outcomes (compared to other possible outcomes) that are deemed important.

Bradley does consider one version that does not fall prey to the problem of gratuitous prevention of goodness (2006: 107). This account, CSSALSC (the Cullity-inspired Self-Sacrifice Absolute Level Satisficing Consequentialism) is dismissed by Bradley—rightly—because it permits self-interested actions that seem obviously impermissible.

5. An alternative suggestion for modifying consequentialism, rather than focusing on effort, is to recast what is required in terms of cost to the agent. This type of modification could also offer a response to demandingness objections. One method would be by applying an absolute restriction to what could be required of an agent. However, this has serious problems. As Ashford notes, sometimes morality may require extreme sacrifices, such as “hacking off your leg with a fairly blunt machete” (2003: 274), so any level must be very high. However, an absolute level of sacrifice this high would not help to avoid the typical demandingness cases, so a fixed restriction on costs to the agent will not help. Alternatively, one might attempt to limit the costs to an agent relative to a given situation, by allowing (but not requiring) agents to give costs to themselves a higher ‘weighting’ in their moral deliberations. This is the proposal made by Scheffler (1982) to allow agent-centred prerogatives, but these have significant issues, famously discussed by Kagan (1989).

on the agent” (2019: 254). If we accept that an action’s difficulty—how much it requires us to apply ourselves as agents—can be relevant to an action’s obligatoriness, this undermines Mulgan’s worry.⁶

We can think of the process for figuring out whether an act is permissible, according to this basic version of effort-satisficing, by going through three steps:

1. Identify what amount of effort can legitimately be required of an agent in this situation.⁷
2. Note how much good the agent can bring about at this level of effort.
3. The agent is then required to bring about *at least* this amount of good. Any action that brings about at least this amount of good is permissible.

As described here, effort-satisficing still permits gratuitous prevention of goodness. To see that this is the case, consider an example from Joe Horton:

Suppose that two children are about to be crushed by a collapsing building. You have three options: do nothing, save one child by allowing your arms to be crushed, or save both children by allowing your arms to be crushed.

(2017: 94)

In this case, Horton supposes—as seems plausible—that the huge sacrifice to the agent makes it the case that they are not required to do anything, that is, it is permissible not to help at all. But if the agent does decide to help, it looks impermissible to save one child, when they could, at no further cost, save both children.

However, if we examine this case from the position of Simple Effort-Satisficing, it suggests that saving one child is permissible. We can see this by going through the aforementioned steps. We have supposed that it is permissible for the agent to decide not to help in this situation. Thus, (as step 1) the amount of effort required is none. The most amount of good they can do (moving onto step 2), while expending no effort, is no good—that is, letting the situation unfold, and leaving both children to die. The agent is then required to bring about at least that amount of good, but any amount above that is permissible. As sacrificing their arms to save one child seems like a better outcome than doing nothing, this must then be permissible. Yet, as those who have discussed such cases have noted (e.g., Horton 2017; Pummer 2016; 2019; Slater 2020), this looks wildly implausible.

6. I do not argue for this claim here, but the moral relevance of difficulty has been addressed in recent papers by people other than Chappell, e.g., Cohen (2000), McElwee (2016; 2022), Nelkin (2016), van Ackeren (2018).

7. Chappell sees this as coming from accepted norms of praise and blame (2019), but to present the general structure, free of special features of Chappell’s account, we can leave this as something to be filled in however is deemed appropriate.

With some tinkering, however, the effort-satisficer can avoid giving this kind of verdict. Chappell's 'willpower satisficing' (WSC) is one example of a more sophisticated effort-satisficing (2019). Chappell requires that when an agent does exert more effort (or use more willpower) than X , they must bring about consequences as good as possible, at that amount of effort. We can characterise this general view as follows:

Efficient Effort-Satisficing: An act φ is permissible iff:

1. φ brings about at least as much good as could be achieved by expending X effort.

AND

2. If the agent expends more than X effort, they must bring about consequences as good as possible, at that level of effort.

The second conjunct guarantees that the agent acts efficiently, that is, that they bring about as much good as they could, given how much effort they are willing to expend. Chappell's WSC is a version of Efficient Effort-Satisficing, where utility is explicitly mentioned (rather than simply "good"). Chappell also supplements his view with a procedure for arriving at the effort ceiling X . On his view, this should be determined by non-consequentialist norms, such as those governing our blame responses.

Now, if we take Horton's case, we can see that saving one child will not be permissible. If, as Horton assumes, doing nothing is permissible, the effort ceiling is at zero, that is, the agent is not required to expend any effort. The second conjunct requires that if any effort above this (above zero) is expended, they must bring about as much good as possible at that level of effort. So saving just one child, when another could be saved at no extra effort, would not be permissible according to this version. Thus, the effort-satisficer can deliver the desired verdicts. This move would be unavailable to the outcome-satisficer because they does not regard effort expended by the individual as a relevant consideration. The efficient effort-satisficer, like Chappell, is able to avoid the charge of gratuitous prevention of goodness, because the second conjunct requires that any prevention of goodness must not be gratuitous.

This appears to give the effort-satisficer a clear advantage. If they are able to give a plausible story about determining how much effort is required in a situation, they will be able to respond to demandingness objections. Where one could act in a very self-sacrificial way, but we do not think this is required, the effort required—which may be physical or motivational—looks like a good reason for why this is the case. They can also avoid the fatal objection for outcome-satisficing consequentialism, as gratuitous prevention of the good is prohibited. For an agent to permissibly choose an option that brings about less good, it must require less effort.

Despite these promising implications, I argue that a revised version of Mulgan’s original objection can still be levelled at the effort-satisficer.

2. Reviving The Problem

Recently, I argued that the type of effort-satisficing described leads to verdicts in many cases that intuitively strike us as too harsh (2020). The types of cases discussed are instances where a person can choose between three options: 1) do nothing, 2) put in a certain amount of effort, bringing about some amount of good, and 3) put in the same effort, but bring about more good. This case structure is identical to Horton’s. Because it condemns gratuitous prevention of the good, effort-satisficing forbids the second option in these cases. While this looks appealing in cases like Horton’s, in other instances, it condemns behaviour that seems permissible, like taking your mother to your favourite restaurant, rather than hers (when doing nothing is also permissible), even though the total happiness would be *slightly* higher if you went to the latter.⁸

While these cases do identify a class of counterintuitive verdicts for the effort-satisficer, someone committed to the view may not find them compelling. They could, for instance, accept that in those sorts of instances, the agent does something wrong, but that it is a minor wrong, and perhaps not one that warrants criticism. Even considering my previous objection, the efficient effort-satisficing view still fares better than maximising act consequentialism, in terms of dealing with demandingness objections *in general*, as it permits agents to act imperfectly.

The typical objection against the outcome-satisficer—condoning gratuitous prevention of goodness—is a much more devastating objection. This is particularly salient in the murder cases. With this in mind, I offer a modification of the original style of case, which I hold commits the effort-satisficer to an extremely unpalatable conclusion.

8. A particularly controversial subset of this type of case that I did not discuss are instances that are entirely self-regarding. Imagine that you could stay on the couch—at no willpower cost—or walk to a vending machine and purchase any of selection of drinks, which cost the same amount. Assuming that there are no morally relevant differences in terms of how this may affect others (e.g., the drinks are not owned by companies which would do different things with the money, the packaging is just as recyclable), willpower satisficing is committed to the claim that you would do something morally wrong by knowingly choosing a drink that does not bring you the most good. Picking a can of lemonade if a cola drink would make me happier would be morally wrong. This implication is seen as so implausible by Urmson that he described it as “preposterous” and “fit for little more than the halting eristic of philosophical infants” (1976: 130). There are of course ways consequentialists can defend these implications (e.g., Singer and de Lazari-Radek 2016: 197), but attributing wrongness in these cases does seem a little odd.

Before I describe such a case, it is worth noting just how disastrous this kind of result is for the satisficer. It has long been a complaint for ordinary (maximising) consequentialists that they allow (or indeed, *require*) murders, as with the surgeon who can secretly assassinate one innocent bystander to provide organs for five of their patients (e.g., Thomson 1976: 206). This offers a challenge to consequentialists, but one for which they have various readily available responses. For such a verdict, they can at least reply that the murder, as horrific to our senses as it might seem, at least leads to the best consequences. They can claim that our intuitions are confused about such a case, because murders *usually* have terrible consequences. In the problem cases for the outcome-satisficer, however, this response is not available. Not only are they committed to condoning acts that strike us as intuitively appalling, but they defend this even when there was a better option available *by their own lights*. This is why this type of case is so problematic for the outcome-satisficer, and why it is such a significant issue if effort-satisficing views also deliver such verdicts.

Recall Mulgan's trolley case, where an agent can do nothing, shoot Bob—the innocent bystander—prompting his corpse to halt a trolley heading off a cliff, or push a heavy sandbag.⁹ If we assume that pushing the sandbag and shooting Bob are equally difficult—where the difficulty includes psychological difficulty—then shooting Bob would be impermissible. However, if shooting Bob is actually easier, but still above the minimum amount of effort required in this situation, this would be permissible. Once again, the satisficer would be committed to permitting murder.

To delve further, we can give a simple 'recipe' for three option problem cases (like Horton's). An agent has three options. All options are permissible—because no effort is required in the situation and each option is efficient. First, they may do nothing. Second, they may expend some amount of energy (say five effort points) to generate some good (say 10 hedons).¹⁰ Third, they may expend *slightly* more energy, and in doing so bring about *much, much* better consequences (say 10,000 hedons). Efficient Effort-Satisficing entails that the second option is permissible. If the second option, however, strikes us as obviously wrong—perhaps because it involves murdering Bob, when one could expend *slightly* more effort to push a heavy sandbag, saving everyone—this is a strong candidate for a counterexample.

An easy response to the amended murder case is that the agent could be required, in this kind of situation, to put in the effort required to push the sand-

9. Mulgan's full case actually includes five options, but I have no need for the complications offered by other options here.

10. I realise that speaking in terms of 'effort points' and 'hedons' is very artificial, but I think it can help to make use of quantities in this kind of situation. It seems like we can make sense of the claim that one act is slightly harder than another, or that it has much better consequences, and this is all these quantities are intended to show.

bag, that is, that however the effort-ceiling (X) is determined, it should require putting in the effort to push the sandbag in this case. As the best outcome they can cause at that level of effort is saving everyone (and not murdering anyone), they would be required to do that.

In Mulgan's case, this looks like the right way to go. Moreover, this is probably the result we would get from applying WSC. This is because Chappell's effort ceiling (X) is fixed by the norms governing our reactive attitudes. As we would blame someone unwilling to expend the amount of effort required to push the sandbag, an agent would be required to bring about as much good as one could by expending that much effort, that is, they could be required to push the heavy sandbag.

However, we can follow the 'recipe' to devise cases within this structure so that it really doesn't look like one is required to do anything. Consider again Horton's example. Because the sacrifice of one's arms is such a significant loss, it does look plausible that the agent doesn't have a requirement to act in this case.¹¹ Now, let us amend the case, so that saving both children requires not only that you lose both of your arms, but has some feature that makes it *very slightly* more difficult. Perhaps, it would involve touching the second child and catching their cold, and, because you know this, it poses a very small added obstacle in your decision-making. *Surely*, in a case like this, if you sacrifice your arms, you would have to save both children. Saving one child in this instance, strikes me as obviously morally impermissible. Efficient Effort-Satisficing does not give this result. Importantly, so long as it is accepted that one would not be blameworthy for failing to act in this case (because reactive attitudes determine the location of the effort ceiling), it provides a counterexample to Chappell's WSC.

We can manufacture many such cases. The second option, which is slightly easier than the best option, could include all sorts of terrible behaviours, including murders. Any theory that permits murder when one could—with the smallest possible amount of extra effort—bring about a better consequence that doesn't involve murder, looks radically implausible.

Before I discuss some ways to revise this view, I would like to address a potential concern with the argument I offer here. One might suggest that by bringing in considerations like the apparent importance of avoiding murder, I beg the question against the consequentialist. Does my argument rely on intuitions a consequentialist would reject, and does this then make these intuitions inadmissible in making a case against the consequentialist (of any flavour)? I do not think so. We can, and should, evaluate consequentialism using non-consequentialist notions. It was evaluation of this sort—intuitions of overdemandingness—which motivated the move

11. If you are a reader still unconvinced, you may ramp up the cost that an agent would suffer, so that they might die rather than losing both arms, or even die a very painful death. Adjust as necessary so it appears plausible that the agent has no moral obligation to save the children.

to satisficing in the first place. Chappell rightly regards it as a virtue of his account that it yields intuitive results. If, as I have argued, it also yields some counterintuitive verdicts, this is a cost to the theory. Just as the defenders of extremely demanding forms of consequentialism have argued that our intuitions about demandingness are mistaken, Chappell could do the same in the cases I have provided. This does nothing, however, to detract from the legitimacy of the complaint.

3. Revising Effort-Satisficing

Even Chappell thinks the types of cases I have described would be a significant problem for his view. He says that “it’s never permissible to do just a little good when a huge amount of good could be achieved by an only slightly more effortful action” (2019: 255). But cases that fit the recipe I described indicate that his account suggests such acts *are* permissible.

A good response on Chappell’s behalf should avoid this implication. One might think this can be achieved with a minor adjustment to Efficient Effort-Satisficing. One possible revision takes the following form:

Revised Efficient Effort-Satisficing (REES): An act φ is permissible iff:

1. φ brings about at least as much effort as could be achieved by expending X effort.

AND

2. If the agent expends more than X effort, they must bring about consequences as good as possible, at the **ballpark level of effort**.

This does seem like a huge improvement. Consider how this revision would operate in practice. An agent must bring about at least the amount of good that could be brought about by expending X effort.¹² If they decide to put in *more* than X effort, they must evaluate whether they could have done more good in their ‘effort ballpark’. If, for instance, they realise they could do more good with slightly more effort, they are required to put in the required additional effort.

This would preclude shooting Bob, when pushing a heavy sandbag (at slightly more effort) would obviously more good. It would also prohibit the second option in the modified Horton cases described above.

Furthermore, it seems plausible that our intuitions might track something like the content of the second condition. We are not sensitive to extremely fine-grained distinctions in effort or willpower. Even introspectively, determining

12. Again, I refrain from any commitments about how this should be calculated, in order that these comments may apply generally, rather than specifically to Chappell’s account.

whether one action or another would require more willpower is difficult. When we are evaluating the actions of others, this is even more difficult, due to different people finding certain acts more cognitively, physically or motivationally difficult.¹³ Very small differences in effort might be indiscernible to us. So perhaps our propensity to blame might depend on whether more good could have been done at the rough level of effort exerted.

However, I contend that REES leads us to a further dilemma. Under this account, once an agent has stipulated some amount of effort to see whether their action is permissible, they are required to ‘scan’ proximate effort levels. If at one of these proximate effort levels, they could do *more* good, they are required to increase their effort levels accordingly. At *this* point, must they again scan the ballpark? Here the proponent of the account faces a dilemma.

On one hand, they may say no. We limit the account to one ‘upwards- revision’ of effort. This option will give some strange verdicts. Consider a case where a given level of effort, *Y*, is such that (however we choose to understand the ballpark level), any increase of effort *within* the ballpark returns *slightly* more good. This account must require an agent considering expending effort *Y* to expend the largest amount of effort in the *Y* ballpark, say *Y+*. However, it may be the case that *slightly* more than *Y+*, the lowest effort possible in the *Z*-ballpark, would result in significantly more good, that is, that there are small marginal returns *within* an effort ballpark, but significant returns *between* ballparks.¹⁴

To describe this in figures that might emphasise the point, perhaps an agent views contributing £50–£99 as requiring roughly the same effort level,¹⁵ but would wince at the prospect of parting with £100. If £50 would save a child, and every penny after that can add some tiny benefit, REES would condemn donating the £50, as £99 would be better, and is within the same effort ballpark. But a £100 donation would be in a different effort ballpark.¹⁶ So REES may endorse donat-

13. For a careful consideration of interpersonal differences in difficulty, see Bradford (2015: 26–63).

14. Of course, there is a vagueness about this talk of approximations, but if we allow that we can distinguish between different rough levels of effort, there will be problems that ensue at the boundary levels.

15. I am using a financial example here for simplicity. This is in-keeping with Chappell’s view, as he accepts that donating money requires exerting willpower (e.g., 2019: 253), and we can easily think of it taking more effort to persuade ourselves to give higher amounts. The same results could be reached with examples using physical effort, but I think we can easily imagine donating within different broad ranges taking more effort, or, in McElwee’s terms, being more ‘motivationally difficult’ for us (2017).

16. It might seem as though units of effort are being conflated with units of money here. I do intend this to track effort, but this coheres with our psychology. Some units of currency do strike us as particularly weighty, so parting with them may disproportionately be motivationally difficult. For example, I still find the prospect of paying more than £5 for a pint in a bar outrageous (and to make offering to buy drinks more difficult!), despite having accepted previous similar price rises without having really noticed.

ing £99 to save a child and buy them nice toys/outfits (or confer some other small benefits), despite the fact that £100 would save two children. I take this conclusion to be somewhat absurd, and once again, it would fail on Chappell's own terms.¹⁷

If this account only requires scanning the effort ballpark once, this can result in strange verdicts when there is an elasticity of utility between the perceived effort levels, that is, where the amount of utility produced is very sensitive to changes in effort at the boundary between different 'ballparks'.

Alternatively, we might think REES should require repeated scanning of the effort ballpark. According to this option, I suppose some amount of effort I consider expending, then I consider what good could be done at similar levels. I then locate the most good I could do at those. At this point I would *again* scan similar effort levels, until I no longer reach an 'effort plateau', where I cannot do more good at a slightly higher effort level.

This interpretation of REES can lead to excessive demands, where the amount of good that an agent can do increases steadily with the amount of effort. After every revision they can scan the ballpark, and realise that they can do more, and more. They essentially would move between 'overlapping ballparks'. This is particularly worrisome for the view because this is the structure of extreme demands in ordinary philanthropic cases: every 'little more' you donate has a negligible effect on you, but added together the demand is enormous.¹⁸

An alternative way we might revise effort-satisficing would be not to require the *most efficient* action if one puts in more than X effort, but to require simply that they don't overlook *significantly* better options at that rough level of effort.

Restricted Effort-Satisficing (RES): An act φ is permissible iff:

1. φ brings about at least as much good as could be achieved by expending X effort.

AND

17. One response possibly available to the satisficer here—and in other places where satisficing gives what looks like permissive verdict—would be to claim that a donation like this, because of how little extra effort would lead to a much better outcome, was suberogatory, i.e., bad but not forbidden (Driver 1992). This move might be possible, but would still leave the view (in this instance) requiring a bad (but permissible) action. This result looks like a failing for the account.

18. Consider, for example, Garrett Cullity's generation of the 'Extreme Demand' (2004: 78) or Fiona Woollard's discussion of the extensive demands placed on pregnant women and mothers (2018). I take no stance here on whether enormous demands like this might be required in some actual instances (though I am sympathetic to that claim). If the view I have described as plausible—that one is required to put in slightly more effort when doing so can bring about significantly more good—is correct, then extreme demands like this could be required if the moral benefits of a further action will *always* lead to significantly more good. This would entail that we can at least describe cases with demands that burdensome on an agent, even if such scenarios could never be realised.

2. If the agent expends more than X effort, there must not be an option ψ requiring a similar amount of effort that brings about *significantly* more good than φ .

This option, at first glance, looks like it would give desirable results in many of the cases described so far. We might think that killing Bob would be impermissible because there is an option requiring a similar amount of effort—pushing the sandbag—which produces significantly better consequences. This view of the account would also avoid the problem cases of being too restrictive (e.g., Slater 2020). If one does act supererogatorily, that is, puts in more effort than required, but fails to maximise the good at that value, this would be permissible, so long as there is not a significant difference in the value of the outcomes. In my Mother’s Day example (2020: 114), for instance, going to either restaurant would be permissible, as there would be no requirement to bring about the *best* outcome at the given effort level—just not to fall very far short of the best.

However, there is a big problem. We must consider how we should interpret ‘significantly’. It has two obvious readings we can disambiguate.¹⁹ ‘Significant’ could have a comparative reading, so that “significantly more” would be assessed proportionally to the amount good done (or that could be done), or it could have an absolute reading, such that some amount of good is significant whatever the circumstances might be.

On either of these options, the view faces problems. On the comparative reading, killing Bob (when this is *slightly* easier than pushing the sandbag) would be permissible in some cases. The determining factor would be the numbers of people saved. If killing Bob saved ten people (ten lives saved, one killed), perhaps one death is a significant fraction, but if there were 100 or 1000 on the trolley (an impressively large trolley, of course), it might not be. So once again the effort-satisficer would get away with murder.

On the absolute reading, we can avoid this result. The death of one person can be regarded as significant, *whatever* the situation is. However, on this view the effort-satisficer will deliver extremely demanding verdicts whenever a life can be saved at relatively little effort. Imagine a charity appeals for donations following a natural disaster. A fairly small sum of money can save a life. Imagine you donate a portion of your month’s salary. You then realise that for a few dollars more—which would require a similar amount of effort (you wouldn’t even notice the difference)—you could save another life. Then at *that* level of effort you have the same realisation. This threatens to be as demanding as Singer’s position in “Famine, Affluence, and Morality” (1972), except with the bizarre

19. These are reminiscent of the two types of satisficing consequentialism Thomas Hurka distinguishes (1990).

caveat that it would still be permissible to do the bare minimum (maximising the good that can be achieved while expending X effort). This looks like the worst of both worlds, as it would give a permissive verdict to people doing very little, and a harsh verdict—the very kind of harsh verdict that motivated the move to satisficing consequentialism in the first place—on people who do try to promote the good more than is typically expected of them.

Another option for interpreting “significantly more” would be as a function of the marginal utility (or marginal improvement to the consequences, for non-utilitarian consequentialists) *and* the marginal effort. We might, for instance, divide the marginal utility of an alternative action (in hedons) by the marginal effort cost (in effort points), and if the result is sufficiently high, we could regard this action as doing *significantly* more good.

According to this view, for any candidate action, if there is an alternative that requires very little additional effort, but yields vast amounts of additional good, an agent will be required to opt for the alternative (at least rather than the original). But if the alternative requires a lot of extra effort for a slightly improved outcome, this will not be required. It will not be easy to specify exactly the ratio of marginal willpower to marginal goodness that meets the threshold, but viewing ‘significance’ in this way looks very promising. It will rule out the exact type of case I have raised as a problem, where a *slightly* more difficult action with much better consequences can be overlooked.

As an added bonus, this modification also delivers intuitive verdicts in other instances of non-maximising beneficence. For instance, many consequentialists want to regard it as permissible for an agent not to donate to the most effective charity when there is an alternative that holds a special meaning for them.²⁰ This account may allow for that kind of intuition, depending on how much worse the preferred charity is and how much extra willpower it would take for them to give to the more effective charity. And it may still condemn donating to a *much* worse charity, which seems like a good verdict.

However, this type of account yields problems when dealing with agents with unusual psychologies. To illustrate my point, recall Mulgan’s trolley case, where an agent may choose between shooting Bob to stop the runaway trolley, or pushing a heavy sandbag. In this case, it is not permissible to do nothing, but the case may be modified to change this. Let us instead imagine that the agent is unable to do anything from their current location, but may take a dangerous path to the sandbag, which will result in the loss of her legs (the danger may be of whatever flavour of your choosing—crocodile, booby traps, or whatever!).

20. E.g., Portmore’s position-relative consequentialism which permits sub-optimal actions (agent-neutrally) when those actions are best from the agent’s position (2003).

They may also take an equally dangerous but different²¹ path to the gun, from where she could shoot Bob. If the paths are horrible enough, it seems like it is permissible for the agent to do nothing.

So far, the active options (the sandbag or the gun) are equally difficult. But imagine that our agent is motivationally unusual. Perhaps because of a deep hatred of Bob, or a lifelong dream of shooting someone in a situation like this, she has a very strong urge to go towards the gun. So strong is her desire that taking the route to the sandbag would be painful for her. While for most people, it might be equally difficult to motivate themselves to go to either route (probably *more* difficult to go to the gun, knowing that the purpose would be to kill someone), for our agent, it would take vastly more willpower to head to the sandbag.

If the required willpower is sufficiently high (lowering the result reached by dividing the marginal good by the amount of effort), our amended RES makes it permissible for this agent to kill Bob. This is not a *gratuitous* prevention of goodness, because there is some reason for her, but this still looks like terrible grounds for condoning a murder. The feature of the case that would make it (apparently) permissible is the relative ease of this option, which only manifests due to the hatred our agent has for Bob. Surely this is not an acceptable justification. Again, the satisficer is left condoning murder, and in a case where a non-murder option is readily available with better consequences. Such a result looks embarrassing for consequentialists.²²

There are several ways the satisficer might respond to this challenge. It is available to the effort-satisficer to suggest that the fact that the difference in difficulty was due to a pernicious motive (the desire to kill Bob) is irrelevant. What matters, they may insist, is how psychologically difficult an action is for an agent to perform, not the desirability of the motives that explain this difficulty. I concede that this is a possible response. However, would I regard this as a highly counterintuitive cost for the view. I take it that those attracted to consequentialism will typically balk at the suggestion that a murder may be permissible, when an alternative action (which is *physically* no more difficult) would lead to a much better outcome.

In support of surviving the verdict, the satisficer could note that such a case is extremely rare. Perhaps the type of case is so unusual that our intuitions are ill-suited to judge it, a suggestion consequentialists have made before (e.g.,

21. I have stipulated that this is a different path to ensure the decision is synchronic rather than diachronic, in case this may make any difference.

22. Consequentialists may deny that such a case is actually possible. The burden of proof for that kind of claim lies with the satisficer. Different people have different psychological profiles, and to argue that a certain profile is impossible seems like a difficult task. Assuming such a case is possible, they may accept the verdict, or attempt to further modify the view.

Unger 1996). However, the gratuitous prevention of the good cases proposed by Mulgan are seen to be sufficiently problematic to motivate a move away from traditional outcome-satisficing (despite also being somewhat fantastic). The only addition for this case is an unusual psychological profile. We accept that some people *do* find some actions more difficult than others (e.g., Bradford 2015), and we do not typically see this as a barrier to applying blame judgements. So, I think this type of case does provide a serious challenge.

There are many additional ways the view could be modified to avoid this verdict. I will briefly discuss one; indexing effort-level to that of an ideal or typical agent.²³ One way to avoid condoning the murder would be to change the way willpower is used in the formulation. Rather than referring to the effort required by an *actual* agent, they could evaluate actions based upon how much effort would be exerted by a typical or idealised agent performing the action. Because typical agents do not harbour a strong desire to kill Bob, pushing the sandbag (saving the same number of lives and killing none) would not require significantly more effort, and would yield significantly more good. Thus killing Bob would be regarded as impermissible.

This would escape the problem case I've described, but the modification comes at a cost. One of the reasons to direct our attention towards effort or willpower is, as Chappell puts it, that “demands on our will—our very *agency*—have . . . the distinctive phenomenology associated with normative demands, in addition to matching up extensionally” (2019: 254). Any move to typical or ideal agents would lose this motivation. Such a move would thus undermine one of the motivations for accepting any form of effort-satisficing.

Once again, there is no easy escape from damaging implications. One way or another, the satisficer is committed to abandoning some of the features that initially made the view appealing.

4. Concluding Remarks

One of Slote's stated goals in proposing satisficing consequentialism was to reconcile consequentialism with common sense moral thinking. His broad strategy—outcome-satisficing consequentialism—has not been a live option as a moral theory since the problems of gratuitous prevention of the good were

23. Other options include revising the axiology such that features of actions we find abhorrent are regarded as having a significant disvalue, regardless of the additional consequences. An example of this would be Amartya Sen's 'goal rights system' (1982). Another option would be “consequentializing”—incorporating some agent centered norms into an agent's evaluations (e.g., Dreier 2011). These strategies represent such significant digressions from typical consequentialist views that discussing their merits here would mean going down a rabbit-hole.

popularised. Before Chappell's contribution, effort-satisficing consequentialism had not been seriously explored as a way to rectify these issues, but given the plausibility of the claim that effort is a morally relevant consideration, the option looks promising. Because of the link between difficulty and cost to an agent, effort-satisficing may be able to accommodate some of the intuitions that led to accounts that explicitly adjusted for the sacrifice of an agent, like Scheffler's agent-centered prerogatives (1982). Exposing the problems that such a theory faces is thus an important project. I have shown here that while the sophisticated version of effort-satisficing consequentialism (like WSC) can avoid the traditional objection, it is susceptible to a slightly revised version of it. This objection looks like a very serious problem for the effort-satisficer.

I have also considered two ways that effort-satisficing might be revised in light of the criticisms offered. I find each of these to have serious problems. In each of the problem cases, a problem arises because the consequentialist ranking of actions does not map nicely onto ordinary deontic evaluations. Specifically, the options that are second-best from a consequentialist perspective (i.e., which have the second-best consequences) but which contain a murder strike us as forbidden, but the worst option from a consequentialist perspective (doing nothing) strikes us as permissible. Whenever a style of case can be given which exhibits this structure, the satisficer's verdict will appear counterintuitive.

Thus, the satisficer is posed with a dilemma. They must accept counterintuitive verdicts, such as allowing murders even when a better outcome was possible, or attempt to further modify the account such that it delivers the verdict that these actions are wrong. Chappell's attempt at the latter (WSC) succeeded in ruling out some of these actions, but I have demonstrated that similar cases can still be generated.

Acknowledgements

This paper has been significantly improved from earlier versions as a result of fruitful suggestions from many colleagues and friends. In particular, I would like to thank Brian McElwee, Lewis Ross, Justin Snedegar and Emilia Wilson. I have also benefited from feedback by University of Glasgow students, and several anonymous reviewers.

References

Ashford, Elizabeth (2003). The Demandingness of Scanlon's Contractualism. *Ethics*, 113(2), 273–302.

- Bradford, Gwen (2015). *Achievement*. Oxford University Press.
- Bradley, Ben (2006). Against Satisficing Consequentialism. *Utilitas*, 18(2), 97–108.
- Chappell, Richard Yetter (2019). Willpower Satisficing. *Nou[^]s*, 53(2), 251–63.
- Cohen, G. A. (2000). If You're an Egalitarian, How Come You're So Rich? *The Journal of Ethics*, 4(1/2), 1–26.
- Cullity, Garrett (2004). *The Moral Demands of Affluence*. Clarendon Press.
- Dreier, Jamie (2011). In Defense of Consequentializing. In Mark Timmons (Ed.), *Oxford Studies in Normative Ethics*, (Vol. 1, 97–119). Oxford University Press.
- Driver, Julia (1992). The Suberogatory. *Australasian Journal of Philosophy*, 70(3), 286–95.
- Horton, Joe (2017). The All or Nothing Problem. *The Journal of Philosophy*, 114, 94–104.
- Hurka, Thomas (1990). Two Kinds of Satisficing. *Philosophical Studies*, 59(1), 107–11.
- Jamieson, Dale and Robert Elliot (2009). Progressive Consequentialism. *Philosophical Perspectives*, 23, 241–51.
- Jenkins, Carrie and Daniel Nolan (2010). Maximising, Satisficing and Context. *Nou[^]s*, 44(3), 451–68.
- Kagan, Shelly (1989). *The Limits of Morality*. Oxford University Press.
- McElwee, Brian (2016). What is Demandingness? In M. van Ackeren and M. Kuhler (Eds.), *The Limits of Moral Obligation: Moral Demandingness and Ought Implies Can* (19–35). Routledge.
- McElwee, Brian (2017). Demandingness Objections in Ethics. *Philosophical Quarterly*, 67, 84–105.
- McElwee, Brian (2022). Cost and Psychological Difficulty: Two Aspects of Demandingness. *Australasian Journal of Philosophy*. Advance online publication. <https://doi.org/10.1080/00048402.2022.2042574>
- Mulgan, Tim (1993). Slote's Satisficing Consequentialism. *Ratio*, 6(2), 121–34.
- Mulgan, Tim (2001a). How Satisficers Get Away with Murder. *International Journal of Philosophical Studies*, 9(1), 41–46.
- Mulgan, Tim (2001b). *The Demands of Consequentialism*. Clarendon Press.
- Nelkin, Dana Kay (2016). Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness. *Nou[^]s*, 50(2), 356–78.
- Portmore, Douglas (2003). Position-Relative Consequentialism, Agent-Centered Options, and Supererogation. *Ethics*, 113(2), 303–32.
- Pummer, Theron (2016). Whether and Where to Give. *Philosophy & Public Affairs*, 44(1), 77–95.
- Pummer, Theron (2019). All or Nothing, but If Not All, Next Best or Nothing. *Journal of Philosophy*, 116(5), 278–91.
- Scheffler, Samuel (1982). *The Rejection of Consequentialism*. Oxford University Press.
- Sen, Amartya (1982). Rights and Agency. *Philosophy & Public Affairs*, 11(1), 3–39.
- Singer, Peter (1972). Famine, Affluence, and Morality. *Philosophy & Public Affairs*, 1(3), 229–43.
- Singer, Peter and Katarzyna de Lazari-Radek (2016). Doing our Best for Hedonistic Utilitarianism. *Ethics & Politics*, 18(1), 187–207.
- Slater, Joe (2020). Satisficing Consequentialism Still Doesn't Satisfy. *Utilitas*, 32(1), 108–17.
- Slote, Michael (1984). Satisficing Consequentialism. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 58, 139–63.

- Thomson, Judith Jarvis (1976). Killing, Letting Die, and the Trolley Problem. *The Monist*, 59(2), 204–17.
- Unger, Peter (1996). *Living High and Letting Die*. Oxford University Press.
- Urmson, J. O. (1976). The Interpretation of the Moral Philosophy of J. S. Mill. In Philippa Foot (Ed.), *Theories of Ethics*, 128–36. Oxford University Press.
- van Ackeren, Marcel (2018). How Morality Becomes Demanding Cost vs. Difficulty and Restriction. *International Journal of Philosophical Studies*, 26(3), 315–34.
- Woollard, Fiona (2018). Motherhood and Mistakes about Defeasible Duties to Benefit. *Philosophical and Phenomenological Research*, 97(1), 126–49.