Consequentialist Options

JUSSI SUIKKANEN
University of Birmingham

According to traditional forms of act-consequentialism, an action is right if and only if no other action in the given circumstances would have better consequences. It has been argued that this view does not leave us enough freedom to choose between actions which we intuitively think are morally permissible but not required options. In the first half of this article, I will explain why the previous consequentialist responses to this objection are less than satisfactory. I will then attempt to show that agents have more options on consequentialist grounds than the traditional forms of act-consequentialism acknowledged. This is because having a choice between many permissible options can itself have value.

I. INTRODUCTION

Act-consequentialism is a combination of two elements.[1] The first of these elements is *axiological*. In any given situation, there is a set of actions which an agent could do. Call these her options. According to the axiological element of act-consequentialism, the agent's options can always be ranked in terms of how much value their consequences would have.[2] Here the consequences of the actions include also the doings of the actions themselves.

The second element of act-consequentialism is *deontic*. It states that an action is right if and only if it is ranked first in the evaluative ranking of the actions which an agent could do in the given situation. If there are many actions which are all ranked first, then the agent is free to do any one of them. This entails that, according to act-consequentialism, all other actions are wrong. Act-consequentialism therefore requires agents to bring about as much value as possible.

The aim of this article is to defend consequentialism against a classic objection. According to this objection, act-consequentialism

[1] S. Sheffler, *The Rejection of Consequentialism* (Oxford, 1982), p. 1.

[2] Many consequentialists tend to rank the options in terms of how much *agent-neutral* value their consequences would have (Scheffler, *Rejection*, p. 1; P. Pettit, 'The Consequentialist Perspective', in M. Barron, P. Pettit and M. Slote, *Three Methods of Ethics* (Oxford, 1996), pp. 92–174, at 130–1). This entails that the same outcomes are ranked in the same order relative to all agents. Others rank options in terms of how good their consequences are relative to agents (D. Portmore, 'Position-Relative Consequentialism', *Ethics* 113 (2003), pp. 303–32). This creates different rankings of the same outcomes for different agents. I take no stand on whether agent-neutral or agent-relative value is to be maximized. Likewise, I take no stand on whether actual or expected value of consequences should be taken into account.

© Cambridge University Press 2014
Utilitas Vol. 26, No. 3, September 2014
doi:10.1017/S0953820814000053

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

fails to provide sufficient *moral freedom*.[3] In the next section, I will describe this objection. In section III, I will turn to why the previous consequentialist responses to the freedom objection are less than satisfactory. After this, in section IV, I will begin constructing my own positive response to the objection. This section will describe how consequentialists can argue that being able to choose an option from a set of many morally permissible options can make the consequences of that option better. The final, and fifth, section will then use this idea to develop a new consequentialist reply to the freedom objection.

II. MORAL OPTIONS AND THE FREEDOM OBJECTION

Jill has just eaten breakfast. She is undecided about what to do next. She could return to bed, watch television, prepare for the next week, visit her elderly parents, assist her local homeless shelter and so on. Most of us think that Jill would not do anything morally wrong by choosing any one of these options. This would be true especially if none of the listed options would have any exceptional consequences. Of course, not all actions which Jill could do would be morally permissible. If she happened to choose to kill her neighbours, she would certainly act wrongly.

According to act-consequentialism, any one of Jill's options would be a morally permissible one only if none of the other alternatives would have better consequences. *All* the above-mentioned options could then be morally permissible alternatives for her only if they all had equally good consequences and there were no other options that would have better consequences.

It is safe to say that this would not be true in Jill's situation. It is often thought that how good the consequences of an option are depends on how much the option brings about pleasure, knowledge, well-being, significant relationships and so on. Consider then Jill's seemingly morally permissible options. It seems unlikely that they would bring about exactly the same amount of value.[4] That various options of different types would each maximize the good seems highly unlikely.

[3] S. Kagan, *The Limits of Morality* (Oxford, 1989), pp. 2–3; Pettit, 'The Consequentialist Perspective', pp. 163–9; Portmore, 'Position-Relative', p. 306; M. Slote, 'Satisficing Consequentialism', *Proceedings of the Aristotelian Society, Suppl.* 58 (1982), pp. 139–64, at 143–9; P. Vallentyne, 'Against Maximizing Act-Consequentialism', *Contemporary Debates in Moral Theory*, ed. J. Dreier (Oxford, 2006), pp. 21–37, at 26–7.

[4] I thus assume that any plausible theory of value will be complete and fine-grained. This is to assume that the evaluative properties of the consequences will not be incommensurable and that the consequences of different options are rarely equally good (Vallentyne, 'Against', p. 26). Whether rejecting this assumption would help act-consequentialism to avoid the freedom-objection will be discussed in the sect. III.3 below.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

This means that, according to act-consequentialism, Jill would be required to choose one of the options listed above, namely the one that has the best consequences. If she chose to do anything else, she would be acting wrongly. Many people believe that we should reject act-consequentialism, if it has this awkward normative implication.

According to this *freedom objection*, ordinary agents have normally many *morally permissible options* (that is, *agent-centred prerogatives*). Furthermore, at least some of these morally permissible alternatives have worse consequences than the other alternatives in the relevant situations. Morally permissible options are thus actions that belong to the set of both optimal and suboptimal actions which we are permitted to do. The idea is that a moral theory which does not leave room for a sufficient number of such options would be too restrictive and therefore implausible.

We can distinguish between the following three distinct objections that can be understood as different forms of the freedom objection:[5]

(O1): A theory is false if it is unable to account for the fact that, in most situations, agents have many morally permissible options.
(O2): A theory is false if it is unable to account for the fact that some suboptimal actions are morally permissible.
(O3): A theory is false if the only way it can account for the fact that, in most typical situations, agents have many morally permissible options is by appealing to some false (or unsatisfactory) theory of value.

During the previous informal presentation of the freedom objection, I raised objections O1 and O2. The thought was that many traditional forms of maximizing act-consequentialism must be false because they fail to account for both the fact that in most cases we have a large number of morally permissible options and the fact that some of these options are suboptimal.

In the following, I will consider the ways in which different forms of consequentialism have tried to avoid these two objections. My objections to these responses will not always, however, be that they fail as responses to O1 and O2. I will also argue some of these responses either (i) fail to avoid O3, or (ii) conflict with our particular intuitions

[5] In addition, we could consider many other related objections such as (O4): a theory is false if it is unable to account for morally permissible options which are moral options either to act so as to make things better overall but worse for oneself or to act so as to make things better for oneself but worse overall (D. Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford, 2011), p. 237). However, using such a theoretical standard that assumes a certain sophisticated philosophical view about which acts are permissible to evaluate other ethical theories seems to beg the important questions.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

about which acts are permissible in particular cases. O3 could also be raised against my own solution presented in section V. However, I hope that the theory of value on which it is based (see section IV) is plausible enough to avoid this problem.

Before I proceed, I want to set aside an objection to act-consequentialism which is related to the freedom objection but which should not be confused with it.[6] Let us return to the previous example. It might well be that, if Jill donated all her money to Oxfam, this would have better consequences than any of her other options. This would mean that act-consequentialism would not allow Jill to choose any other option. In fact, she could never do anything else except to use her life to help the people who live in extreme poverty. This seems to require too much from Jill. This objection to act-consequentialism is called the *demandingness objection*.

Even if this objection too illustrates how act-consequentialism constrains too much what an agent is allowed to do, I believe that Peter Vallentyne is right in arguing that the demandingness objection raises a different set of issues than the freedom objection.[7] This is because a moral theory can be demanding even if it leaves a lot of freedom to choose between different actions. It could, for instance, allow us to use any one of the infinitely many ways of helping others. Even if we had this many options, the theory would still require us to give up our own personal projects. This shows that a view may allow us to choose between many different actions which all are demanding in their own way.

Conversely, there could be a moral theory that only required us to stay in bed. This theory would not be very demanding but it would still fail to leave room for freedom to choose between different alternatives. Because of these two possibilities, a satisfactory response to the freedom objection does not necessarily help with the demandingness objection and vice versa. For this reason, I will concentrate here only on the freedom objection.

III. THE PREVIOUS RESPONSES TO THE FREEDOM OBJECTION

Consequentialists have been aware of the freedom objection for a long time and thus it has been discussed extensively. The previous act-consequentialist responses to it fall into three categories. The responses

[6] Mulgan and Pettit run these objections together (T. Mulgan, *Future People* (Oxford, 2006), pp. 17–20, and Pettit, 'The Consequentialist Perspective', pp. 163–9). Vallentyne correctly warns against this (Vallentyne, 'Against', pp. 26–7; see also Scheffler, *Rejection*, p. 10).

[7] Vallentyne, 'Against', pp. 26–7.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

of the first category rely on rejecting our common-sense intuitions about what we are permitted to do. The responses of the second category try to avoid the objection by reformulating the normative element of act-consequentialism. The responses of the third category concentrate on the axiological element of act-consequentialism.

III.1. Resisting intuitions

Some act-consequentialists insist that the freedom objection fails because it is based on mistaken common-sense intuitions. On this view, it only seems to us that we usually have many morally permissible options some of which are also suboptimal. However, this is only an illusion. In fact there is always something which we are required to do. In this way, these consequentialists reject the desiderata for ethical theories on which O1 and O2 are based.

The act-consequentialists who give this response recognize that most of the seemingly permissible options do not bring about the most valuable outcomes. They then argue that it would be 'profoundly irrational' to bring about a suboptimal outcome.[8] This seems to entail that 'there can never be reason to produce less of a value rather than more' and for this reason it would also be wrong to do so.[9] The suboptimal actions which we tend to take to be morally permissible are thus argued to be actually morally impermissible.

There are two ways in which this response can be made more appealing. First, it can be argued that our moral intuitions should be disregarded unless they can be vindicated by a defensible moral theory. Act-consequentialists can then try to show that any moral theory which incorporated moral options could not be given coherent justification.[10] If that were the case, then they would be entitled to claim that we should ignore our intuitions about moral options because there cannot be good theoretical support for them.

Act-consequentialists can also argue that there is a difference between the permissibility of actions and whether we should criticize

[8] P. Pettit, 'Satisficing Consequentialism', *Proceedings of the Aristotelian Society, Suppl.* 58 (1982), pp. 165–76, at 174.

[9] Pettit, 'The Consequentialist Perspective', p. 128.

[10] According to Shelly Kagan, the defenders of many morally permissible options must both (i) accept that there are options to allow harm to happen, and yet (ii) deny that there are moral options to harm others. Kagan then claims that any defence of the morally permissible options to allow harm to happen (usually based on the costs to the agents from not being able to allow harm to happen for others) will also justify options to harm others which the defender of moral options has already rejected. As a result, the views which try to defend ordinary moral options will be incoherent unless the defender of those options can also defend moral constraints against causing harm to others (Kagan, *Limits*, pp. 19–24 and ch. 7). The problem then is that such constraints also seem equally paradoxical (Kagan, *Limits*, pp. 24–32, Scheffler, *Rejection*, ch. 4). For a response to this argument, see Vallentyne, 'Against', p. 24.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

agents for doing them.[11] Some intuitively permissible suboptimal actions could be claimed to be actually morally wrong actions for which the agents are *not* condemnable.[12] It could then be argued that there is nothing problematic about the fact that we should call most of the actions which we could do in any given situation wrong. After all, the wrongness of these actions does not imply a negative evaluation of us as agents. Wrongness of actions would on this view be cheap and harmless.

I have no objections to these responses. If my positive proposal below is correct, then it turns out that they are merely redundant. Consequentialists themselves would have in that case been wrong about how much moral freedom their view provides. If that were the case, there would just be less consequentialist motivation for the previous evasive strategies.

III.2. Reformulating the deontic element

The responses to the freedom objection in the second category try to make room for morally permissible options which include suboptimal alternatives by revising the deontic element of act-consequentialism. In this way, they tackle both O1 and O2 directly.

Even the move from act-consequentialism to *rule-consequentialism* could be understood as this kind of a response. Rule-consequentialists claim that whether an action is right is determined by whether the action would be permitted by the optimific set of moral *rules*.[13] There are many potential sets of moral principles which we could adopt. Whichever set we internalized, this would have different consequences. This means that potential sets of principles can be ranked in terms of how good the consequences of their general internalization would be. According to rule-consequentialism, our actions are then right if and only if they are authorized by the code which would have the best consequences.

[11] For a recent investigation of this difference, see T. M. Scanlon, *Moral Dimensions* (Cambridge, Mass., 2008), ch. 1. Consequentialists also tend to understand criticizing an agent as an action. They then often think that we should do this action when it maximizes the good. It may then be that criticizing someone for failing to maximize the good doesn't itself maximize the good in many situations. For these consequentialists it is natural to think that wrong and that which ought to be criticized come apart. Other consequentialists distinguish between wrongness and appropriateness of reactive attitudes such as blame. These consequentialists can think that it is appropriate to blame even if criticizing an action would not maximize the good. The consequentialists who use the evasive strategy can use either one of these distinctions.

[12] Pettit, 'The Consequentialist Perspective', pp. 164–8.

[13] For sophisticated formulations of the view, see R. Brandt, *A Theory of Good and the Right* (Oxford, 1979); B. Hooker, *Ideal Code, Real World* (Oxford, 2000); Mulgan, *Future People*, chs. 5–9.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

This allows rule-consequentialists to argue that the adoption of moral codes which leave room for agents to choose between many different permissible actions in normal circumstances would have better consequences than the adoption of the more restrictive codes.[14] In this way, rule-consequentialism will make many suboptimal actions too permissible.

For example, rule-consequentialists can claim that having freedom to choose between different permissible actions is a necessary precondition for autonomous agency. They can also argue that autonomous agency is both instrumentally valuable (it brings about more well-being) and intrinsically valuable (i.e. good in itself). Thus, a rule that allows agents to choose between many morally permissible options will have the best consequences because it will bring about valuable autonomous agency.

The second response of the 'normative type' grants agents an explicit permission to pursue some suboptimal outcomes and thus also increases the number of permissible options which individuals have in normal circumstances. This way to avoid O1 and O2 is known as *the hybrid view*.[15] It accepts the act-consequentialist thesis according to which it is always permissible to bring about the best outcome. However, on this view, we *also* in some cases have an additional permission to do some actions which will not have optimal consequences.

When we assess whether a given suboptimal action is permissible, we should take into account two considerations.[16] First, we should consider how much less value this suboptimal action would bring about. Second, we should also take into account how much the agent would need to sacrifice her personal projects, relationships and interests, if she chose to do the action with the best consequences. An agent is then permitted to do either the optimal action or the given suboptimal action, when both (i) the suboptimal action would not have significantly worse consequences than the optimal action and (ii) the optimal action would require significant personal sacrifices. It can be argued that such permissions to do some suboptimal actions must be granted to all agents in order to protect their 'integrity' and to recognize the natural independence of the agents' individual points of views from which they pursue their personal projects.[17]

Both of the previous responses to the freedom objection contain an important insight. Act-consequentialists too should recognize the

[14] Mulgan, *Future People*, p. 165.
[15] Scheffler, *Rejection*, ch. 2.
[16] Scheffler, *Rejection*, p. 20.
[17] Scheffler, *Rejection*, ch. 3.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

value of autonomy, personal projects and commitments, integrity and individual perspectives. I will try to show below how this is possible within a more purely consequentialist framework. I will also argue below that, if we truly recognize these values within that framework, this will generate plenty of moral options for us. If my proposal turns out to be successful, then at least the freedom objection does not give us reasons to move to either rule-consequentialism or the hybrid view.

There is also a third, less radical way of revising the deontic element of act-consequentialism in response to the freedom objection. This is illustrated by the so-called *satisficing act-consequentialism*.[18] According to this view, right actions have 'good enough' consequences. One is permitted to do any action such that either (i) the value of its consequences does not fall below a certain antecedently specified threshold, or (ii) the value of its consequences is reasonably close to the value of the best consequences. Satisficers then argue that most of our intuitively permissible options (including many suboptimal ones) satisfy these criteria.[19]

Unfortunately, if satisficing act-consequentialism were true, then certain intuitively morally wrong actions would turn out to be morally permissible.[20] In some cases, agents will be able to bring about enough value either by harming others or by actively preventing others from getting some benefits. Under the satisficing views, some such actions would be permissible because the value of their consequences will be close enough to the value of the best outcome which the agent could bring about. If we believe that such actions should not be morally permissible options for agents, then we cannot accept satisficing act-consequentialism as a response to the freedom objection.

III.3. Reformulating the axiological element

The third kind of act-consequentialist responses to the freedom objection focuses on the axiological element of act-consequentialism. My own response to the freedom objection will also use this strategy. However, in this sub-section, I want to argue first that the previous attempts to use this strategy have not been successful.

I have assumed so far that the axiological elements of act-consequentialism are both complete and fine-grained. If this assumption were true, then evaluative ties between the outcomes of different options and cases where their goodness is incommensurable would be rare. An act-consequentialist could try to create morally

[18] Slote, 'Satisficing'; Vallentyne, 'Against'.

[19] Slote, 'Satisficing', p. 158.

[20] T. Mulgan, 'Slote's Satisficing Consequentialism', *Ratio* 6 (1993), pp. 121–34; B. Bradley, 'Against Satisficing Consequentialism', *Utilitas* 18 (2006), pp. 97–108.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

permissible options by rejecting my assumption.[21] If we could not compare the value of different outcomes or we could do so only very roughly, then there would often be many actions such that no other actions would have determinately better consequences.[22] Either the consequences of the other actions would be equally good in the rough comparisons or their value would be incomparable. Act-consequentialism would then entail that agents would be permitted to do any one of these actions. In this way, this response directly answers O1. And, if it were right, then there would be no reason to worry about O2. The intuitively permissible suboptimal options could always be argued to have evaluative properties that are incomparable.

Unfortunately, this strategy fails to save all intuitively morally permissible options. A defender of this view has to accept that there are several evaluative properties that are incomparable. For each one of these properties, there will be a distinct set of more basic natural properties in virtue of which things have that evaluative property. The current proposal fails to deal with cases in which the only difference between the options which an agent has is how much they have just one of those basic natural good-making properties.

It seems plausible that the consequences of actions have one sort of value in virtue of how much pleasure they contain. Consider then a case in which Ann's friend Ben visits her. They would both enjoy watching television, especially given that both *The Wire* and *Breaking Bad* are on at the same time. Ann asks Ben which show they should watch. It turns out that both Ben and Ann are indifferent between the two options. In this situation, intuitively it seems permissible for Ann to choose to watch either one of the shows with Ben. However, let us assume that both Ann and Ben would ultimately get slightly more pleasure out of watching *The Wire*. They would experience nine units of pleasure if they watched *Breaking Bad* and twelve units if they watched *The Wire*, and there are no other differences between the outcomes of these actions.

In order to say that it is permissible for Ann and Ben to watch either one of the shows, the act-consequentialist view under consideration must claim that neither one of the two possible outcomes would be better than the other. This could not be because the value of the two outcomes is incomparable. After all, only the value of pleasure is in question, and it must be comparable how much of that single value the outcomes have. So, the defender of this response would have to

[21] Consequentialists who reject the strong commensurability of values include J. Griffin, *Well-Being* (Oxford, 1986) and R. Chang, 'Introduction', *Incommensurability, Incomparability, and Practical Reason*, ed. R. Chang (Cambridge, Mass., 1997), pp. 1–34.

[22] Vallentyne, 'Against', p. 26.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

claim that the outcomes have an equal amount of the relevant single evaluative property.

However, this is not plausible. If pleasure is something which makes outcomes of actions good, then the additional three units of pleasure just cannot fail to make *The Wire* outcome better than the *Breaking Bad* outcome. For this reason, a coarse-grained and incomplete axiology is unable to save all the intuitive moral options. This means that, even if this response could respond to O1 and O2, it cannot generate all intuitively morally permissible options if we assume any plausible theory of value. It thus fails to avoid a version of O3.

Another version of act-consequentialism tries to accommodate the intuitively morally permissible options by relying on many evaluative rankings of outcomes.[23] According to Portmore's *dual-ranking* proposal, from an agent's perspective, the actions which she could do in her situation can be ranked in terms of both *moral* and *all-things-considered* (moral and non-moral) value of their consequences.

For example, saving a child from a burning building whilst risking one's life and failing to get to an important meeting can bring about a state of affairs that is morally speaking the best from one's perspective and yet not the best, all things considered, from one's perspective. This is because the potential losses in one's own well-being do not make the resulting states of affairs morally worse from one's perspective and yet these losses do make the same consequences worse all things considered from that perspective.

Portmore then claims that an action is permissible if and only if it does not have an alternative which is better from one's perspective in terms of both moral and all-things-considered value.[24] This explains why it is permissible both to save the child in the previous example and not to do so. These actions do not have an alternative that would have more of both moral and all-things-considered value from the agent's perspective.

Portmore's proposal faces a problem when we consider trivial actions that are suboptimal on both accounts.[25] Some of these actions are permissible options. Recall the example in which Ann could either watch *The Wire* or *Breaking Bad* with Ben. She enjoys both shows even if she too gets a bit more pleasure out of *The Wire*. In this situation, from her perspective, it would be morally better to watch *The Wire* because as a result of that choice *Ben* would experience more pleasure. The same option would also be better all things considered from her perspective given that it would give both her and Ben a bit more pleasure.

[23] Portmore, 'Position-Relative', sect. 5.
[24] Portmore, 'Position-Relative', p. 322.
[25] Portmore has recognized this himself (Portmore, *Commonsense*, p. 135 n. 22).

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

In this situation, Portmore's dual-ranking view would not permit Ann to watch *Breaking Bad* with Ben instead of *The Wire*. This is because she would have an option that would bring about both more moral and more all-things-considered value. However, it seems that, intuitively, watching either one of these shows should be morally permissible. Therefore, Portmore's view seems to have at least some unintuitive consequences.

Because of this, Portmore has revised his view. For reasons of space, I cannot explain his highly sophisticated view here in full. His basic idea is that a given alternative is a permissible option for an agent if and only if the agent does not have an alternative which she has both more requiring reason and more reason, all things considered, to perform, where a requiring reason is just a reason that has some morally requiring strength.[26]

This new view can easily explain why Ann and Ben are permitted to watch either *Breaking Bad* or *The Wire*. Portmore can plausibly claim that they have no morally requiring reasons for watching *The Wire* but perhaps only so-called 'enticing reasons'.[27] Portmore's new theory can thus deal with the problem I presented for his previous view. The conjunctive condition on the right hand-side of his proposal is satisfied by both actions in my case and thus they both are permissible options for Ann and Ben.

Portmore's new theory thus has the required resources for capturing the extension of morally permissible actions correctly. This is because Portmore helps himself to a distinction between morally *requiring* reasons and the reasons that lack such distinct deontic force. However, this solution has a cost. It was an advantage of his earlier view that it used purely evaluative rankings of outcomes to determine what is required and what is merely permissible.

In contrast, Portmore's new theory takes the basic deontic notion of requiring force of reasons as fundamental. It then explains what is required and what is not required but merely permissible in terms of this fundamental notion. Thus, on his view, what is morally required and what is morally permissible is not a consequence of purely evaluative rankings of states of affairs, but rather a consequence of which reasons can require and which ones cannot.

The alternative I will offer below has at least one advantage over Portmore's view in this respect.[28] It does not take deontic notions

[26] Portmore, *Commonsense*, pp. 137 and 232–3.

[27] J. Dancy, 'Enticing Reasons', *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, ed. J. Wallace et al. (Oxford, 2004), pp. 91–118.

[28] The objection here is not simply that Portmore's view takes reasons to be more basic than value. If the so-called buck-passing theories are correct, then this is true in any case

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

such as the requiring force of reasons as fundamental. In true consequentialist fashion, it takes evaluative rankings of states of affairs to be fundamental and then attempts to make sense of deontic notions such as requirements and permissions in terms of these more basic rankings. In this sense, my view is more explanatory than Portmore's view. Rather than assuming basic deontic distinctions such as a distinct requiring force of certain reasons, it gives an account of requirements in terms of value.[29]

IV. THE VALUE OF HAVING MANY PERMISSIBLE OPTIONS

This section introduces a hypothesis according to which how much value the consequences of an option contain can be affected by whether the agent was required to choose the option or whether she chose it out of many permissible options. I want to suggest that there are ways in which the fact that an agent is allowed to choose an option out of many permissible options can make the consequences of that option better. Perhaps there are also cases in which giving an agent many permissible options can make the consequences of an option worse.[30]

Gerald Dworkin has argued that having freedom to choose an option out of many permissible options can make the consequences of the option better in three different ways.[31] Such choices can have instrumental value, intrinsic value and constitutive value.

An option chosen from many permissible options has instrumental value whenever the chosen option brings about, as a result of having been chosen from many permissible options, other things which are

(T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass., 1998), ch. 2). Rather, the problem is that he has to introduce two new irreducibly deontic dimensions on which the strengths of reasons vary. These are the requiring and non-requiring strengths of reasons. It is no surprise that we can give an account of what we are permitted to do in terms of which considerations require acts from us. My aim below is to avoid relying on such deontic notions in order to give an account of what is permissible.

[29] I believe that we should therefore prefer my view because it can explain more of the deontic realm in terms of the evaluative and without a commitment to additional deontic distinctions in doing so. It is open for Portmore to argue that the deontic distinctions he relies on are independently motivated and so can be made use of without extra ontological cost. In any case, it seems like a prima facie advantage of my view that it can do without such distinctions.

[30] For instance, making a choice between many permissible options can itself have information-gathering and deliberation costs (G. Dworkin, *The Theory and Practice of Autonomy* (Cambridge, 1988), p. 66). For other potentially negative effects of having freedom to choose in certain contexts, see Dworkin, *Autonomy*, pp. 67–78. Carter argues that such negative effects of freedom do not outweigh the positive value of having a choice between many permissible options (I. Carter, 'The Independent Value of Freedom', *Ethics* 105 (1995), sect. 6).

[31] Dworkin, *Autonomy*, pp. 78 and 80.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

intrinsically good. It might be difficult to imagine how a free choice between many permissible options could bring about other goods which the chosen option could not bring about itself without that choice. However, there is some reason to believe that this is possible.

First, it could be claimed that individuals want to be able to choose between many permissible options and that they find making such choices pleasant.[32] In this case, having a choice between many permissible options can be instrumental for bringing about desire satisfaction and pleasure, which are intrinsically good.[33]

Being able to choose an option from many permissible options can also provide us with knowledge about ourselves that we could not learn by any other means. As Dworkin points out, if one wants to know whether one is courageous or cowardly, one can do so only by seeing which one of the permissible options one opts for in certain situations of risk.[34] If one were always required to choose specific options in the relevant kind of risk-scenarios, then there would be no room for the type of practical deliberation in those situations that is required for acquiring self-knowledge about one's character. One could then argue that this kind of self-knowledge is intrinsically valuable, and that options can bring about this good only when they are permissible but not required options for the agents in the relevant situations.

It could also be argued that the fact that an option can be chosen from many permissible options can have intrinsic value as such. Being free to choose an option out of many permissible options would in this case be desirable for its own sake. If this were true, then many options would bring about additional value when one can choose to do them from many permissible options. This is because one consequence of these options in this situation would be that an intrinsically valuable choice out of many permissible options takes place.

Unfortunately, assessing whether the choices between many permissible options really have such intrinsic value is rather difficult (sect. V.5). At best, we can support this idea by the observation that we tend to prefer to get new permissible options even when we do not end up taking advantage of them.[35] In any case, I will not rely

[32] J. Rawls, *A Theory of Justice* (Cambridge, Mass., 1971), p. 143; Dworkin, *Autonomy*, p. 79.

[33] Carter, 'Independent', p. 832.

[34] Dworkin, *Autonomy*, p. 79.

[35] T. Hurka, 'Why Value Autonomy?', *Social Theory and Practice* 13 (1988), pp. 361–82; A. Sen, 'Freedom of Choice: Concept and Content', *European Economic Review* 32 (1988), pp. 269–94; Carter, 'Independent', pp. 830–1.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

on the controversial view according to which choices between many permissible options have intrinsic value.[36]

Perhaps the most important kind of value which the choices between many permissible options can have is the so-called constitutive value.[37] This value does not reside in the causal consequences of options or in the choices between many permissible options themselves. Rather, the idea of constitutive value of choices is that choices between many permissible options can in part constitute larger, more complex wholes which have intrinsic value. More specifically, free choices seem to have two different kinds of constitutive value.

Some free choices between many permissible options have constitutive value which is often called 'representative value'.[38] This type of value can be illustrated with an example of a poem. If a poet writes a poem, the value of the outcome does not solely depend on the poem's aesthetic qualities or on how much the audience appreciates it. The same poem could have been randomly generated by a computer. Yet, the outcome that the given poem exists in that scenario would not be equally good.

The poem gets additional significance from the fact that the poet chose the words out of many permissible options to express her beliefs, desires and emotions. The poem, as a result of the prior choices between many permissible options, comes to represent the poet's thoughts. As such a representation, the poem and the choices which generated it have constitutive value. After all, the poem and those choices are now a part of a larger, more complex whole which includes the poet's thoughts expressed by the poem, the world which these thoughts represent and the audience's appreciation of the poem as a representation of the world. It can then be claimed that this complex whole has intrinsic value.

That the poet is able to choose her words out of many different permissible options thus makes the world better in its own way. This is one example of the constitutive value of choices between many permissible options. This mechanism works in the same way when we think about being permitted to choose gifts from many permissible options for the people we love. These gifts also get additional value from the fact that they express our emotions as a result of the free choice which we made between many permissible options.

[36] The view that choices between many permissible options are intrinsically valuable is criticized in Dworkin, *Autonomy*, p. 80; W. Kymlicka, 'Liberalism and Communitarianism', *Canadian Journal of Philosophy* 18 (1988), pp. 181–204, at 187; R. Dworkin, 'Why We Do Not Have a Right to Liberty', *Liberty and the Rule of Law*, ed. R. L. Cunningham (College Station, 1979), pp. 167–81, at 170–1. It is defended in Carter, 'Independent', sects. 2–3; Hurka, 'Why Value'; Sen, 'Freedom of Choice'.

[37] Dworkin, *Autonomy*, p. 80.

[38] Scanlon, *What We Owe*, p. 252.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

Such choices can also have constitutive value which is sometimes called 'symbolic value'.[39] We intuitively take some outcomes to be better when they result from our free choices between many permissible options. This is perhaps easiest to see in the case of one's own actions. I value my career, friends and projects more because I have been able to choose them freely out of many permissible options. I also value other people's personal convictions and achievements much more when they are consequences of free choices between many permissible options. These things deserve far less admiration when they have been chosen for the agents by others. Thus, free choices between many permissible options can make the world a better place.

Why would this be? Compare the situation in which we are free to make important life-choices between many permissible options to a situation in which we only have one alternative available for us because others have already made the choice for us. When our ability to choose between many permissible options is taken away from us, this can be understood as a judgement that we cannot make reasonable choices ourselves. Such judgements make the resulting outcomes worse from our perspective.

In contrast, granting someone a freedom to choose between many permissible options is a way of recognizing her cognitive capacities to make decisions. It can then be argued that such recognition of our rational faculties is an intrinsically valuable complex whole, and that having a choice between many permissible options is a constitutive part of this whole.[40] This is another reason why choices between many permissible options have constitutive value (called 'symbolic value').

In the following, I will assume that free choices between many permissible options can have value in these ways and perhaps also in some other ways. I will also assume that the amount of value which a free choice between many permissible options has is a very context-dependent issue. The value of a given choice will depend, for example, on what the other permissible options happen to be.

V. CONSEQUENTIALIST OPTIONS

I can then finally turn to my own consequentialist solution to the freedom objection. This proposal will be based on the thought explained in the previous section; the consequences of an option can be better when the agent is able to choose freely the option from a set of many

[39] Dworkin, *Autonomy*, p. 80; Scanlon, *What We Owe*, p. 253.

[40] It could also be argued that choices between many permissible options are a constitutive element of other intrinsically valuable complexes such as human agency and self-respect (Carter, 'Independent', p. 839).

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

morally permissible options. I will first give a more theoretical sketch of this proposal. I will then describe a concrete example which illustrates it. After this, I will deal with two potential counterexamples to my suggestion, explore one more metaphysical objection to it and finally discuss what motivation there is for accepting it. I will then conclude by making one further axiological observation about the new framework.

V.1. The Proposal

My proposal begins from the thesis that, *on purely consequentialist grounds*, a suboptimal option is a morally permissible option for an agent when it is better by consequentialist lights that the agent is morally permitted to choose this suboptimal option, whichever permissible option she ends up choosing. More precisely, on consequentialist grounds, a suboptimal option A is a morally permissible option for an agent whenever A is such that, even if the agent ends up choosing that option when it is a morally permissible option in a large set of permissible options, the outcome will be better than if she chose some other morally permissible option without A being a permissible option for her. After all, in this situation, that the agent is able to choose A as one of her morally permissible options increases the amount of value there is in the world.

This account could be formalized, for instance, in the following way.[41] Assume that, in the given circumstances, action-tokens A_1, or A_2, or . . ., or A_x are an agent's options.[42] Each one of these options would have some definitive consequences if the agent chose it when there was a requirement for her to do so. Let us assume that A_1 would have the best consequences if the agent were required to choose this option, A_2 would have the second-best consequences if the agent were required to choose it, . . . and A_x would have the worst consequences if the agent were required to choose it.

Let us then use curly brackets for alternative sets of morally permissible options which an agent could have.[43] So, if the agent is free to choose between A_1 and A_2 her set of morally permissible options will be {1, 2}, if the agent is free to choose between A_1, A_2 and A_3 her set of morally permissible options will be {1, 2, 3} and so on. My consequentialist proposal for which actions are morally permissible options then is:

[41] I thank Daniel Elstein for this formalization.

[42] Note that the relevant options here include only actions which it is possible for the agent to do. For this reason, whether it is permissible for an agent to just fly away from the situation does not arise.

[43] This element of my proposal raises interesting questions in moral metaphysics. I will discuss these questions in sect V.5.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

Consequentialist options: An option A_n is a morally permissible option just in case the value of the consequences of A_n given the set of morally permissible options {1, . . ., n} is at least as great as the value of the consequences of A_{n-1} given the set of morally permissible options {1, . . ., n-1} where A_{n-1} is the least optimal option in {1, . . ., n-1}.

V.2. An Illustration

In order to illustrate how this proposal works, we need a concrete example. Let us assume that table 1 represents Jill's options in her circumstances. It also represents the value of the consequences of those options both when she is required to choose them and when she can freely choose them out of certain sets of morally permissible options. Here I am assuming that these options are mutually exclusive. If Jill does one of these actions, she will not be able to do any of the other alternatives.

Table 1

	Action	Value of consequences when required	Value of consequences of A_φ given the option-set {1,. . .,φ}
A_1	Charity work	200u	200u
A_2	Go to Paris	150u	210u
A_3	Visit parents	100u	215u
. . .	. . .	. . .	. . .
A_{n-1}	Stay in bed	20u	240u
A_n	Ignore Joe's call	−30u	100u

In this situation, the traditional interpretation of act-consequentialism would require Jill to do charity work. This is because that option has the best consequences (200u) when we compare all options as required ones.[44] However, when we ask whether Jill has, by consequentialist lights, also a morally permissible option to go Paris, we can apply the consequentialist options principle. According to this proposal, Jill has this additional morally permissible option if the value of (a) her going to Paris when she can choose between doing so and charity work is more than the value of (b) the charity work in a situation in which this is required from Jill.

According to table 1, the value of the consequences of (a) is 210u. It is assumed, at this point, thus that the consequences of going to Paris when chosen from the set of morally permissible options of

[44] The value of the consequences of the acts when they are required is just the standard value of their consequences when the freedom to choose the given alternative from some set of morally permissible options is not taken into account. Thus, I assume here that having a set of morally permissible options of just one action cannot make the consequences of the given action better in any way.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

Paris and charity work are 60u better than they would have been otherwise. Adding this value is justified because having free choices between many permissible options adds value to the world (sect. IV). Thus, the consequences of Jill choosing to go to Paris given the set of morally permissible options {charity work, going to Paris} are better than the consequences of her doing charity work when she would have been required to do so (200u).[45] This entails that, according to the consequentialist options principle, it is permissible for Jill to go to Paris too in the described situation.

At this point, Jill's permissible options are either to do charity work or to go to Paris. Does Jill also have a third morally permissible option to visit her parents? In order to answer this question, we should again consult the consequentialist options principle. Table 1 shows that, if Jill visits her parents when it is permissible for her either to do so, go to Paris or do charity work, the value of the consequences will be 215u. According to the consequentialist options principle, we should then compare this amount of value to the value of Jill going to Paris when she would have only the slightly smaller set of morally permissible options of either going to Paris or doing charity work.[46] The value of the consequences of that action in those circumstances is 210u.

This means that, if Jill also has the third option to visit her parents, this will have better consequences than if Jill chose the least valuable options of the permissible options already granted to her. This is why Jill has, on consequentialist grounds, also a third morally permissible

[45] To be clear, here we compare the consequences of three options. Call option A Jill's act of going to Paris when it is permissible for her either to do charity work or to go to Paris. Call option B Jill's act of going to Paris when she is required to do so. Call action C Jill's act of doing charity work when she is required to do so. The claim here is that because of the value of having a free choice the consequences of A have more value than the consequences of either B or C.

This claim might sound implausible. How could whatever little pleasure Jill gets from being able to choose to go to Paris from a set of morally permissible options outweigh the amount of pleasure added by the charity work? This objection relies on the assumption that the defender of the Consequentialist Options principle accepts a hedonistic value theory. But, given that the defenders of this principle can recognize other symbolic and constitutive values discussed above, there's no reason why the defenders of the principle are committed to such an axiology. They can argue that free choices have good-making properties that are more important than the value of pleasure (see sect. V.4 below). The defenders of the principle can also still accept that, even if going to Paris and doing charity work are both permissible options, it would still be better if Jill did charity work.

[46] Here we compare the value of freely visiting parents to whichever free choice would have the least valuable consequences from the smaller set of morally permissible options {do charity work, go to Paris}. I assume that going to Paris would have fewer valuable consequences in this case than doing charity work. This requires that the value of the consequences of doing charity work increases when Jill gets to choose between charity work and going to Paris. Table 1 would need an additional column to illustrate this.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

option to visit her parents. And, by iterating this process, we can step by step determine which morally permissible options Jill has.

According to table 1, Jill would no longer have the morally permissible option of ignoring her friend Joe. All the way up to that point, adding a new option to Jill's set of morally permissible options could add value to the world, no matter which option she chose from the larger set of morally permissible options. But the basic consequences of this option are already negative in value (−30u), and there is not enough additional value in having the freedom to choose this option. If Jill were able freely to choose to ignore Joe, the consequences of this action would be 100u in total. In comparison if she chose to stay in bed when she has the smaller set of morally permissible options in which this is the best option, the consequences would be better, namely 240u.

In this situation, the consequentialist options principle does not give Jill the additional option to ignore Joe. It would therefore only permit her to choose between staying in bed and all the other morally permissible options she would have in the largest set of morally permissible options arrived at by applying the consequentialist options principle that does not contain the option of ignoring Joe. I hope that this brief illustration shows how the consequentialist options principle applies in an ordinary case.

V.3. Consequentialist options and the freedom objection

Of course, the proposal sketched above leaves many details of the view to be addressed. In this section, I have room to discuss only some of them. First of all, we should recognize that the consequentialist options principle itself is merely a formal principle for determining which options are morally permissible on consequentialist grounds. The previous example also illustrates how that principle on its own cannot tell us which options are morally permissible. The facts about the permissibility of different options depend on how much more valuable (if at all) the consequences of options are when the agent is able freely to choose them from different-sized sets of morally permissible options.

Section IV tried to explain certain considerations which can, on many occasions, make the consequences of freely chosen options more valuable. However, already the previous simple example shows how complex are the issues we face here. I merely selected, for illustrative purposes, the values of the consequences of different options so as to get the right intuitively permissible options for Jill. If this were an actual choice-situation, in order to fill in table 1 correctly, we would need to know (i) whether Jill wants to make choices between many permissible options and enjoys making them, (ii) what kind of representational value her particular choices could have (for instance, how much more her *free* choice to visit her parents would mean to them) and (iii) which

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

options Jill must have for her cognitive capacities to be recognized. We would need to have all this information to know which morally permissible options Jill would have according to the consequentialist options principle.

However, the point of introducing that principle is more structural than practical.[47] According to the freedom objection, we tend to have more morally permissible options than act-consequentialism seems to entail. However, if the consequentialist options principle is true, then how many morally permissible options we have depends on how much the value of the consequences of our options is affected by what options we are permitted to choose. If choices between many morally permissible options can make the consequences of options a lot better, then we will have many morally permissible options. In contrast, if no additional value is generated by the choices between many permissible options, then consequentialism would still not create many options.

All of this means that, whatever options we intuitively think we have in any given situation, there will be a version of consequentialism that entails that very number of morally permissible options.[48] There will be a way of assigning additional value for the consequences of options chosen from a given set of morally permissible options such that (i) the consequences of the agent having that set of morally permissible options are not worse than the agent doing the least valuable action from the more limited sets of morally permissible options and (ii) giving the agent a more extended set of permissible options could have worse consequences.

If this is right, then there cannot be an objection to consequentialism *per se* that it does not provide enough permissible options. This means that the freedom objection debate will then no longer be about consequentialism as such, but rather about whether the kind of axiological views that generate the right amount of permissible options together with the consequentialist options principle are defensible more generally. If the required view about the value of having many permissible options does not turn out to be plausible, then one could argue that at least the most plausible versions of consequentialism do not provide all the permissible options we think we have. However, at this point, the consequentialist could claim that she has been able to

[47] I assume that consequentialism should be understood as a criterion of permissibility rather than as a decision-procedure (J. S. Mill, *Utilitarianism* (London, 1861), ch. 2, para. 19). For this reason, it is not an objection to my proposal that the consequentialist options principle will be difficult to apply in practice.

[48] This would mean that, contrary to what Campbell Brown has argued, it is possible to consequentialize moral theories which contain moral options. This is because my version of consequentialism does not entail his 'dominance' condition (C. Brown, 'Consequentialize This', *Ethics* 121 (2011), pp. 749–71).

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

make her view compatible with most of our intuitions about morally permissible options. This would mean that she could now more easily claim that only *some* – but not most – of our common-sense intuitions about options are mistaken (see sect. III.1).[49]

V.4. Counterexamples

At this point, it could be protested that the consequentialist options principle has clearly counter-intuitive implications in some hypothetical cases. For reasons of space, I cannot here pre-empt all potential objections, but I want to explain how the ideas of the previous section help us deal with at least two problematic cases.

Suppose first that Jack really wants to be able to choose between a set of permissible options that includes him torturing a baby for the sake of his own mild amusement. Suppose also that Jack finds having such an option in his option-set intensely pleasurable. It could then be objected that, in this situation, the consequentialist options principle entails that torturing the baby for fun is a permissible option for Jack.

The important thing to observe is that this awkward conclusion does not follow from the consequentialist options principle alone. Rather, it is generated by an axiological assumption according to which Jack's act of torturing a baby has better consequences when this is a permissible option for him because having this option satisfies Jack's desires and gives him pleasure. There is no reason why a defender of the Consequentialist Principle would need to accept this assumption.

The defenders of that principle can accept that satisfying a desire to have an option and experiencing pleasure from having an option have some modest amount of value. This value perhaps makes some alternatives permissible, but not alternatives such as boiling babies which have a lot of bad consequences anyway.[50] It can also be argued that being able freely to choose this particular alternative has little if any symbolic or constitutive value.

[49] This would allow the consequentialist also to reply to one understanding of the so-called integrity objection (see B. Williams, 'A Critique of Utilitarianism', in J. J. C. Smart and B. Williams, *Utilitarianism: For and Against* (Cambridge, 1973), pp. 77–150, at 116–17, and Scheffler, *Rejection*, pp. 7–9). If consequentialism provides a significant number of permissible options to choose from in usual circumstances, then it no longer requires each agent in all cases to produce the best available outcome. Hence, there would not usually be a requirement to neglect one's personal projects. If this is right, then consequentialism will no longer alienate one's integrity in an objectionable way.

[50] In giving this response, I am not assuming any lexical orderings of different values or doing/allowing distinctions. I am merely assuming that a single short pleasure or satisfaction of a single desire has little value when compared to the value of a whole human life.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

This response suggests a general strategy for dealing with counterexamples. Whenever an objection is made that the Consequentialist Principle has counter-intuitive consequences in individual cases, the defender of this principle will consider the axiological assumptions that ground such objections. She will then revise the value which was ascribed to being able freely to choose a given alternative in the relevant context. The ways in which having free choices can be valuable are flexible enough for this strategy.

There are cases, however, in which the defenders of the consequentialist options principle will dig their heels in. Suppose an evil demon wants me to be able to choose between a set of permissible options that includes torturing my baby for the sake of my own mild amusement. The demon will wreak untold havoc on billions of people's lives unless it is permissible for me to do this act. In this case, the consequentialist options principle would imply that it is permissible for me to act in this way.

This is a conclusion which the defender of the consequentialist options principle should endorse. Consider a case in which an evil demon would wreak untold havoc on billions of people's lives unless I tortured my baby for fun. In this case, my intuition is that I should torture my baby for fun.[51] If intuitively I can be required to torture my baby for fun in order to avoid untold suffering of billions of people, then it should be no less intuitive to accept that in some cases for this reason it should be merely permissible for me to torture my baby for fun.

Furthermore, the defender of the consequentialist options principle can still say many moral things about the permissible option in this case. Even if this act would be permissible, it would still be evil, bad, what I have good reason not to do, wicked and so on. Here the point is that we can make many moral and normative evaluations of the morally permissible options that an agent has that have nothing to do with the permissibility of these options. As a theory of permissibility, the consequentialist options principle in itself is neutral about these other moral qualities of actions.

V.5. A metaphysical objection to the view

Let us then consider a more metaphysical objection to my proposal. Most ethicists agree that, if a certain action-token has a certain moral properties such as permissibility, then it has those moral properties necessarily in all identical situations. The consequentialist options principle has to agree with this. Once we have used it to determine which options are permissible in a given situation, these options are

[51] Hooker, *Ideal Code*, pp. 129–31.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

necessarily permissible. They are permissible in all otherwise identical worlds.

However, when we apply the consequentialist options principle itself, we are required to consider otherwise identical hypothetical circumstances in which agents have different sets of morally permissible options available for them. Strictly speaking these situations are then impossible states of affairs. I do not believe that this is a problem as such for the view.

This is because it is commonplace to talk about such situations in ethics. Philosophers often begin from an ethical principle such as 'coercion and deception are always wrong' to which their opponent is committed.[52] They then consider which actions are permissible under this principle and what the consequences of those actions' permissibility would be. In these cases too, ethicists typically describe and discuss impossible states of affairs in which the deontic facts are different from what they are necessarily. The application of the consequentialist options principle relies on this very same type of hypothetical situations that are strictly speaking metaphysically impossible but at the same time at least epistemically possible. In metaphysics, there is a number of appealing views about the metaphysical status of such impossible situations.[53]

The opponent of my principle could argue that, even if the metaphysically impossible situations are metaphysically acceptable, there is no way of determining how much value such impossibilities would contain. I admit that it is very difficult to evaluate how much value worlds that contain square circles would have. Despite this, I believe that this epistemological worry is exaggerated.

First, we are only required to evaluate situations that are otherwise exactly like our own actual situations except that different actions are morally permissible in them. We are thus not required to evaluate situations that contain square circles or people feeling *Schadenfreude* at the good fortunes of others.

Second, as the example above illustrates, these situations are epistemic possibilities because they are conceivable in a strong sense. This is illustrated by the fact that we can understand moral views with which we do not agree. We know, for example, what it would be like for it to be permissible to kill infidels. In this situation, there would be sufficient reasons to kill them and doing so would be justified, it would not be appropriate to blame anyone for killing infidels, you should not feel guilt for doing so and so on.

[52] D. Parfit, *On What Matters*, vol. 1 (Oxford, 2011), pp. 178–9.

[53] D. Nolan, 'Impossible Worlds: A Modest Approach', *Notre Dame Journal of Formal Logic* 38 (1997), pp. 535–72.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

The final question then is: do we have a method of evaluating how valuable the relevant epistemically possible but metaphysically impossible situations would be? I described in section IV the ways in which having many permissible options could be valuable. This could be instrumentally, intrinsically and constitutively valuable.

Evaluating instrumental and constitutive value of different actions will not pose any special epistemic challenges. We can consider how much pleasure would be experienced in these worlds, how many desires would be satisfied, what agents would know about themselves and how valuable such organic wholes as poems, gifts, relationships and careers (in which choice plays an essential role) would be in the relevant situations. We would do this exactly in the same way as we evaluate the value of these considerations in our own world.

This leaves us with the intrinsic value of different option-sets in the relevant epistemically possible but metaphysically impossible situations. This is one of the reasons why, in formulating my proposal, I did not want to rely on the intrinsic value of having many morally permissible options (sect. IV). For what it is worth, I believe that our own carefully considered convictions can offer a reliable guide to tricky epistemic questions like this. For one, the fact that we tend to prefer having many morally permissible options (we would prefer to live in worlds in which we would have more options) can be accepted as evidence of the fact that having many morally permissible actions has intrinsic value. Perhaps by considering what we would be willing to trade for a larger number of permissible options, we could come to understand how intrinsically valuable morally permissible options are. I admit that this quick response to the epistemic problems of my view will probably not convince everyone.

Because of this, it is worthwhile to recall that we can also work in the other direction. As section V.3 explained, we can also begin from our intuitions about what is permissible and then use the consequentialist options principle to evaluate how much value we give for having many morally permissible options. In this way, my proposal does not even require being able to determine independently how much intrinsic value different sets of morally permissible options would have.

V.6. The arguments for the consequentialist options principle

So far, I have not provided an argument for the consequentialist options principle. I have only argued that, together with certain prima facie plausible axiological assumptions, this principle can be used to construct a consequentialist theory which fits our intuitions about which options are morally permissible.

Of course, this in itself does not give us a decisive reason to accept my proposal. Even if my proposal has one advantage, there might be

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

other forms of consequentialism that have even bigger advantages. Unfortunately, for reasons of space, I cannot offer here a comprehensive comparison of my view and the other forms of consequentialism. I can, however, offer two considerations to motivate my suggestion.

First, my consequentialist proposal is based on the very same attractive general beliefs about morality as many other forms of consequentialism. Like other forms of consequentialism, it too takes evaluative notions to be primary and then explains deontic notions such as obligations and permissions in terms of the evaluative rankings of the states of affairs.[54]

It is also based on the attractive idea that links morality to maximizing the amount of value in the world. Like act-consequentialists, it too can claim that 'The dictates of any alternative view will at least sometimes prescribe "avoidable" misery, or at least missed opportunities for benefit.'[55] According to my proposal, the dictates of act-consequentialism prescribe missed opportunities too because, if act-consequentialism were true, then agents would have fewer opportunities to make free choices between many morally permissible options. As a result, there would be fewer symbolic, constitutive and other forms of value in the world. In this way, my proposal can accommodate the attractive general Kantian ideals of freedom and autonomy within the consequentialist framework.

I can also offer a 'Pareto improvement' argument for why my proposal is more plausible than many traditional forms of act-consequentialism. Either being able freely to choose between many morally permissible options is valuable or it is not. If it is not, then the consequentialist options principle collapses into extensional equivalence with act-consequentialism. In that case, according to my principle, an agent would not have any other morally permissible options than the one option that has the best consequences. As a result, the consequentialist options principle would have the same benefits and problems as the standard forms of maximizing act-consequentialism, and so it would be equally good.

In contrast, if the consequences of options can be better when agents have a larger number of morally permissible options, then it seems like the consequentialist options principle would be more plausible than the standard forms of act-consequentialism. There would then be one form of value which the traditional consequentialist views could not capture whilst describing which outcomes we should promote. Furthermore, in

[54] M. Smith, 'Neutral and Relative Value after Moore', *Ethics* 113 (2003), pp. 576–98.

[55] Hooker, *Ideal Code*, p. 6, see also J. J. C. Smart, 'Outline of a System of Utilitarian Ethics', in J. J. C. Smart and B. Williams, *Utilitarianism: For and Against* (Cambridge, 1971), pp. 3–74, at 5–6.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

this situation, the consequentialist options principle would also fit our moral intuitions about which options are permissible better than the traditional forms of act-consequentialism.

I admit that these arguments do not amount to a decisive argument for my theory. My more modest aim has been to introduce a consequentialist alternative which has at least one important attraction. Arguing for it in more detail must be left for another date.

5.7 One further observation

Let me finish by making one further observation. It concerns the axiology which we would need to generate a sufficient number of morally permissible options. Consider the permissible options ranging from A_1 (charity work) to A_{n-1} (staying in bed) in the earlier example. If we ignore the value of being able to choose from many permissible options, A_1 has the best consequences, A_2 has slightly worse consequences and so on, all the way to the almost neutral consequences of A_{n-1}. Yet, every time Jill is granted a new permissible option (an option less valuable in itself) there will be more value in the world than there would be without that permissible option.

This requires that, when we go down the scale of actions and towards the larger sets of permissible options, at each step more choice-based value must be added to the consequences of options when they are added to the ever-growing sets of permissible options. The difference in value between doing A_2 when that act is required and doing it with the set of permissible options {1, 2} was 60u, the difference between doing A_3 when that act is required and doing it with the set {1, 2, 3} was 115u and so on.

To some extent, such a pattern of increases in the value of choice is plausible. If you have only two options, you would want to have more of them, your actions would not really be representative of your thoughts, your rational capacities would not really be recognized and so on. So, being able to choose the second option out of the set of permissible options containing two options would not add much value to the consequences of that option. However, if you had three permissible options instead, maybe you would get a bit more of the choices you wanted, the results of your actions would be a bit more representative of your inner world and your abilities would be a bit more recognized. Similar increases would presumably take place if we added a fourth permissible option, a fifth one and so on.

However, at some point, the added value of the new permissible options would begin to decrease due to the diminishing marginal utility of having more morally permissible options to choose from. The value added to the consequences of options by the fact that there are more permissible options to choose from would at that point

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press

stop compensating for the fact that the other consequences of these remaining options are always a bit less valuable. It is difficult to assess whether this would happen too quickly for the consequentialist framework to create enough permissible options. Despite this, I conclude that the consequentialist options principle is well worth investigating further as a form of consequentialism that perhaps could avoid the freedom objection.[56]

j.v.suikkanen@bham.ac.uk

[56] I would like to thank Douglas Portmore and Daniel Elstein for their assistance.

https://doi.org/10.1017/S0953820814000053 Published online by Cambridge University Press