




Supererogation and sequence

Adam Bales¹ · Claire Benn² 

Received: 5 February 2019 / Accepted: 22 November 2019 / Published online: 11 February 2020
© The Author(s) 2020

Abstract

Morally supererogatory acts are those that go above and beyond the call of duty. More specifically: they are acts that, on any individual occasion, are good to do and also both permissible to do and permissible to refrain from doing. We challenge the way in which discussions of supererogation typically consider our choices and actions in isolation. Instead we consider sequences of supererogatory acts and omissions and show that some such sequences are themselves problematic. This gives rise to the following puzzle: what problem can we have with a sequences of actions if each individual act or omission is itself permissible? In this paper, we develop a response to this question, by exploring whether solutions analogous to those proposed in the rational choice literature are available in the case of supererogatory sequences. Our investigation leads us to the view that making sense of the supererogatory requires accepting that there are global moral norms that apply to sequences of acts alongside the local moral norms that apply to individual acts.

Keywords Rational choice · Supererogation · Sequences · Permissibility · Money pump · Norms

1 Overview

Actions that go above and beyond the call of duty—supererogatory actions—are common in our everyday lives and thinking. These actions are morally better than

Adam Bales and Claire Benn contributed equally to this paper, and the order of author names was chosen at random.

✉ Adam Bales
atb39@cam.ac.uk

Claire Benn
cmabenn@gmail.com

¹ St John's College, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DA, UK

² The Polonsky Academy, The Van Leer Jerusalem Institute, 43 Jabotinsky Street, Jerusalem 9214116, Israel

other actions that could permissibly have been performed and yet are neither morally required nor morally forbidden.¹ The types of actions that are supererogatory range from the saintly and heroic (risking one's life helping those with Ebola) to the mundane (sending a 'get well' card to a colleague).

Take a classic example:

THE BURNING BUILDING: you encounter someone trapped inside a burning building. After calling the fire department, you enter the building and save the person.

Going into the building was neither morally required nor morally forbidden: it was morally optional. It was also morally good: you saved someone's life. Given that it is morally permissible for you to call the emergency services and do no more, entering the building is morally better than a permissible alternative.² Consequently, entering the building is supererogatory.

The central puzzles in the supererogation literature concern two features of this notion: optionality and goodness. For example, if an action is the morally best action, why aren't we required to perform it?³ Alternatively, how are we to balance our own wellbeing against the wellbeing of others?⁴

In this paper, we set these classic puzzles aside and pose a new one. In particular, we challenge the way in which discussions of supererogation typically consider our choices and actions in isolation.

To get to this challenge, note that the above case could be filled out with additional details that would change our judgement as to whether saving the person from the burning building was supererogatory. For example, if you had to kill someone else to save the person in the building, it might no longer be permissible to do so. Likewise, if the person in the building carries a plague that will kill millions if they are rescued, perhaps it isn't good to save them. Additionally, the optionality and goodness of an act may depend on the actions that an agent could have performed instead (Benn 2017). Now, what these factors have in common is that they relate to the actions and alternatives currently available to the agent, and their future consequences. On classic articulations of supererogation—on what we will call a 'myopic' analysis⁵—whether or not an act is supererogatory depends only on such factors.

However, we will discuss whether a largely-overlooked consideration is also relevant to whether or not an act is supererogatory: the other supererogatory actions you have or have not performed in the past and will or will not perform in the future.

¹ The 'morally better' condition rules out as supererogatory cases where our duty can be fulfilled by either of two (or more) equally good actions. Such actions are optional, given that each is not forbidden (it fulfils a duty after all) and is not required (as the alternative could be performed instead). Yet we would not consider them to be supererogatory: one of them must be performed in order for us to do our duty.

² Entering a burning building is also more costly than merely calling for help. It is therefore supererogatory even on accounts where a supererogatory act must involve a cost to the agent, for example Rawls (1999) and Benn (2017). In any case, in this paper we make the working assumption that optionality and goodness are sufficient for supererogation (though many of our claims will hold for various accounts that deny this assumption).

³ This is the 'Paradox of Supererogation' (Horgan and Timmons 2010; Heyd 1982).

⁴ See Raz (1975).

⁵ We draw this terminology from the rational choice literature. See McClennen (1997).

In the first half of this paper, we introduce three puzzles for supererogation that arise when we consider sequences of choices. In these cases, myopic analyses prove inadequate for explaining our intuitions fully. In the second half of the paper, we show that this challenge is similar to some well-known challenges from the literature on rational choice. By reflecting on the relationship between the two issues, we argue that making sense of the supererogatory requires paying attention to the sequences in which actions are embedded.

2 Three puzzles

We start with three puzzling cases involving supererogation.

2.1 Scrooge and the Saint

First:

SCROOGE: Ebenezer Scrooge, in Charles Dickens' *A Christmas Carol*, is a person of principle. As Scott C. Lowe argues, "he believes that his actions are justified and that others have no right to demand more of him. Indeed, judged by his actions alone, Scrooge is a moral man, if all we mean by that is that he does not violate common moral principles. Nowhere in the story do we read of Scrooge lying, cheating or defrauding anyone" (Lowe 2009, p. 29). Scrooge is a person of principle, in the sense that he scrupulously does his duty, but he refuses on every occasion to do more. He therefore never performs a supererogatory act.

However, despite Scrooge always doing his duty, there is something problematic about the fact that he deliberately does no more than what duty demands of him on each and every occasion.⁶

THE SAINT: The Saint always chooses the supererogatory act; she always chooses the act that is morally better for others no matter the cost to herself. She forgives everyone the wrongs they do her; she spends all of her time working for the good of others; she has no projects of her own or, if she does, when they conflict with other people's needs she gives them up.⁷ Throughout her entire life, she never chooses self-interest over moral good.

Scrooge never goes beyond duty, while the Saint continually puts aside her own needs for the needs of the others. The former, therefore, never performs a supererogatory act, while the latter performs them whenever possible. Interestingly, however, there is a sense of disquiet in both cases. Scrooge is, in some sense, without fault—he never

⁶ Our use of the term 'problematic' is deliberate ambiguous: we do not want to adjudicate here on whether we should view Scrooge's actions (and the other actions to be discussed) as bad or morally prohibited. At this point, we simply want to highlight that there is something of concern morally speaking about Scrooge's behaviour.

⁷ Susan Wolf distinguishes the Loving Saint, whose concern for others plays the role that more selfish concerns play in our lives, from the Rational Saint, who has concerns like ours but who sacrifices his interests for the interests of others (1982). For our purposes, either type of saint will do.

does anything forbidden, never shirks duty—and yet is quintessentially reprehensible. Meanwhile, the Saint seems superhuman in her selflessness, the epitome of moral virtue, and yet there is something worrying about her behaviour. After all, as Jean Hampton notes when discussing a similar case, extreme selflessness can mean a loss of the self.⁸ Ultimately, Hampton argues, “not all self-sacrifice is worthy of our respect or moral commendation” (Hampton 1993, p. 1).

The puzzle is how to account for our concern about both sequences of action, given that it is permissible on any occasion to perform a supererogatory action and also permissible to refrain from doing so. Thus, on each occasion, it is permissible for Scrooge to refuse to go beyond the call of duty, just as it is permissible for the Saint to choose to do so.

So we have our first two puzzles. First, given that each of Scrooge’s omissions is apparently permissible, what can be morally problematic about the sequence whereby he never performs any supererogatory acts? And second, given that each of the Saint’s acts are permissible, what can be morally problematic about the sequence whereby she performs every supererogatory act?

2.2 The colleague

Our first two puzzles involve sequences of actions and omissions that span a lifetime. However, the issue we are interested in does not arise only over a whole life. It can arise for an individual who performs a shorter sequence of actions in a specific context. Consider:

THE COLLEAGUE: Sara has 100 work colleagues. She gives 99 of them presents on their birthdays but decides not to give the 100th a present.

Considered in isolation, Sara’s decision to not buy a gift for just one of her colleague is morally unproblematic: gift-giving is supererogatory and it is thus permissible to refrain from giving gifts. However, the sequence of decisions where Sara buys all but one of her colleagues a present is morally problematic. So, again, a puzzle arises: how do we account for the fact that this sequence is morally problematic when it is made up of decisions that seem morally unproblematic?

These three puzzles present a challenge, one that myopic accounts of supererogation are ill-placed to respond to. After all, these accounts focus solely on the current action, while the puzzles under discussion arise as a result of an action’s place in broader *sequences* of actions.

Further, the challenge deepens once we note that two natural responses fail to respond adequately to these puzzles. According to the first of these, there is nothing problematic about the above sequences of actions themselves. Instead, it might be suggested, what is problematic in each case is the agent’s character. For example, perhaps what is problematic in THE COLLEAGUE is just that Sara displays a thoughtless

⁸ Hampton points to the case of Terry who dedicates her life to her household responsibilities despite a series of heart-breaking events (1993, p. 1).

character.⁹ Indeed, we agree with half of this claim: we find it plausible that Sara displays some defect of character. However, we deny that this is *all* that's problematic in THE COLLEAGUE: in addition to considerations of character, we think it clear that there is something problematic about the actions themselves.¹⁰

Moving to the second response then, it might be denied that all acts in the above cases are really permissible. One way of spelling this out starts by noting that past actions can determine the normative status of present actions. For example, whether or not I have a duty to water your plants while you're on holiday depends on whether I've previously promised to do so. That is to say, my act of promising changes the normative status of refusing to water the plants. Similarly, if two medicines are fatal if administered together, but cure a disease if administered individually, then whether I should administer the second medication depends on whether I administered the first (Portmore 2013). So the act of administering the first medicine changes the normative status of administering the second. Returning to our puzzle cases, a similar line might be taken. For example, it might be suggested that Scrooge acts impermissibly in refusing to donate on later occasions, precisely because of his earlier refusals.¹¹

Ultimately, we think this is a promising response. However, we think it's the start of a conversation, not the end: it's not enough to simply insist that the earlier acts determine the status of the later acts. Instead, we must be given an explanation as to why this would be so. In the case of promising, the earlier act very explicitly involved committing to a later act, so the link between the two is clear. In the case of the medicines, the earlier act changes the (non-normative) consequences of the latter act, so again the link between the two is clear. Yet the link in the above cases is not so apparent, so there is more work to be done. We will return to this issue below, once we have more background to hand. In the meantime, we set this response aside and so take the three puzzles to remain.

3 Resolving the puzzles: first steps

At the core of these puzzles is an incompatibility between our evaluations of sequences and our evaluations of the acts comprising these sequences. This incompatibility results from the plausibility of:

HARMONY: A sequence of acts is morally problematic only if at least one of the acts that comprise the sequence is morally problematic.¹²

⁹ Heyd, for example, claims that such behaviour can “reflects a particular contempt, aversion or mean intention” (1978, p. 31).

¹⁰ Those who maintain that all three cases can be explained by appealing only to considerations of character will be uninterested in our full account. Nevertheless, the discussion to follow will help clarify how moral considerations are influenced by sequential factors. Someone who sees the problems with our sequences as being problems of character can read this as providing an account of how sequential considerations influence evaluations of character (rather than how they influence evaluations of acts).

¹¹ Thanks to a reviewer for pushing us on this point.

¹² HARMONY is similar to several other principles discussed in the rational choice literature. Those interested in exploring these might consider the discussion of the Traditional View in Seidenfeld (1994), the discussion

HARMONY is plausible: it is natural to expect our evaluation of a sequence to parallel our evaluation of the acts that comprise this sequence. However, this principle is what generates our puzzles. Each act in the three sequences discussed, *on its own*, seem like a permissible commission of a supererogatory act or a permissible omission of a supererogatory act. However, if we accept that the sequences as a whole are problematic, then HARMONY entails that this initial analysis must be incorrect. This is where the tension lies. Typically, issues raised by HARMONY are overlooked in discussions of supererogation, which tend to focus on isolated acts rather than sequences.¹³ However, rational-choice theorists have spilled much ink on sequential matters. So, we turn now to three solutions to sequential puzzles developed by rational-choice theorists. With each, our question will be: can an analogous solution resolve our supererogation puzzles?

3.1 The denial solution

In the rational choice literature, a paradigm sequential puzzle is the money pump argument against cyclic preferences. An agent has such preferences if she prefers a first thing to a second, the second to a third, and then the third to the first (hence, her preferences start from the first thing and cycle back to it). Example: an agent who prefers apples to bananas, bananas to carrots, and carrots to apples.

The money pump argument demonstrates that an agent with such preferences can be made to spend money for no benefit. To see how, note that the above agent will be willing to pay a small cost to make any of the following trades:

1. an apple for a carrot (because she prefers carrots to apples);
2. a carrot for a banana (because she prefers bananas to carrots);
3. a banana for an apple (because she prefers apples to bananas).

That is, letting arrows indicate a willingness to pay to move from one item to another: Apple → Carrot → Banana → Apple.¹⁴

Now, imagine that our agent has an apple and we offer her three trades. First, we offer to trade her apple for a carrot, for a small cost (5 cents, say). From (1), above, the agent will take this trade. Then, again for 5 cents, we offer to trade her carrot for a banana. From (2), the agent will take this trade. Finally, for 5 cents, we offer to trade her banana for an apple. From (3), the agent will take this trade. However, she has now paid 15 cents and once more owns the apple she started with. So she has spent money for no gain. While each act (of trading for a preferred fruit) seems rational in

Footnote 12 continued

of the Deal Agglomeration Principle in Arntzenius, Elga, and Hawthorne (2004), and the discussion of the Packaging Principle in Hajek (2009).

HARMONY will be plausible only if we treat two acts as different if they have different non-normative consequences. For example, in the case mentioned in §2.2, giving the second medication will be one act if the first medication was administered and a different act otherwise. The important point is that our puzzle cases apparently continue to violate Harmony, even if we differentiate acts in this way.

¹³ Exceptions include Driver (1992, see footnotes 34 and 35) and those who address supererogation and imperfect duties such as Portmore (2016).

¹⁴ For a prominent case where such preferences seem to arise, see Quinn (1990).

isolation, the sequence of acts seems problematic. So our evaluation of the sequence is incompatible with our evaluations of the acts that comprise it.¹⁵

The standard response to this puzzle is to deny that an agent may rationally have cyclic preferences.¹⁶ Indeed, on this view an agent's acts can't even be assessed for rationality if she has such preferences. Given this, the money pump argument collapses, as it doesn't make sense to discuss whether the sequence of trades is rational: having cyclic preferences renders this question moot.

Does this illuminate our supererogation puzzles? Well, on one construal, the above response denies the phenomena (cyclic preferences) a place in the world of rationality. Likewise, we could deny the phenomenon that gives rise to supererogation puzzles (the supererogatory) a place in the moral world; perhaps no acts are supererogatory, and we are always morally required to carry out the morally best action (call this the *denial solution*).¹⁷

To see how this resolves the puzzles, consider SCROOGE: Scrooge regularly fails to carry out the morally best action and so, given the above, regularly acts immorally. Therefore we can retain HARMONY: Scrooge's sequence of acts is problematic but so too are some of the constituent acts. This restores harmony between our evaluations of the sequence and the acts. No puzzle remains.

Yet this solution is unsatisfying. After all, this paper is addressed to those who take supererogation seriously enough to be concerned by the puzzles outlined above. For many such people, in Brandt's words, an ethical theory without supererogation "can hardly be taken seriously: like the Sermon on the Mount, it is a morality only for saints" (Brandt 1979, p. 276). So, rejecting supererogation should be a last-ditch solution, to be adopted only if less costly solutions are not available. We will argue that other solutions are available, so set aside the denial solution.

3.2 Satan's apple and Chinese takeaway

So we're assuming a prior commitment to the existence of the supererogatory. On the other hand, it's widely felt that the rationality of cyclic preferences can be dismissed with ease. Given this, there is a crucial disanalogy between the supererogation puzzles and the cyclicity puzzle: the puzzling nature of the former must be taken seriously, while the latter can be set aside. So, it's worth finding a closer analogue to the supererogation puzzles.¹⁸ Better, let's find two.

First, consider a puzzle due to Arntzenius, Elga, and Hawthorne:

SATAN'S APPLE: Satan has cut an apple into infinitely-many pieces. Eve will now be offered these pieces, one at a time, at an ever-increasing pace, so that

¹⁵ Perhaps this money-pump argument must be replaced with a more complex variant (see Rabinowicz (2000)). However, the simple argument suffices for our purposes.

¹⁶ See Davidson et al. (1955)

¹⁷ Different approaches can be analogical to some solution in different regards. So, there are multiple solutions to the supererogation puzzles that might be described as analogical to the stated solution to the money pump argument. Here, we will focus on the analogy that we have identified.

¹⁸ The paradox of the preface might also provide an interesting analogy (Makinson 1965). However, we will focus on decision-theoretic comparisons.

she will be offered them all in a finite period. All else being equal, Eve prefers more apple to less. Yet if Eve eats infinitely-many slices, she will be banished from Eden. Eve strongly desires to avoid banishment (Arntzenius, Elga, and Hawthorne 2004).

Here's the puzzle: each time Eve is offered an apple piece, it seems permissible to accept. To see why, start by noting that when Eve is offered a piece of apple, she knows that she will either accept infinitely-many pieces of apple on other occasions or she won't. If she does, she is doomed to be cast from Eden whatever she now does, so she might as well take the extra piece of apple. If she doesn't, she won't be cast from Eden whatever she now does. Again, then, she might as well take the extra piece of apple. So whatever Eve expects to do on other occasions, she should take the apple now. Yet this reasoning applies equally to each piece of apple and so Eve will say yes to each piece, and hence to infinitely-many pieces. So while each act seems permissible, considered individually, the sequence of acts leads to banishment. Again, the sequence seems problematic, despite the comprising acts seeming unproblematic. Further, Eve's preferences seem unimpeachable: she wants apple and fears banishment. There's nothing strange about such preferences. So, this puzzle is harder to dismiss than the cyclic-preferences puzzle.¹⁹

To get to the second analogous case, consider Lydia who is deciding whether to order pizza or Chinese takeaway. She doesn't prefer one to the other and still wouldn't prefer one to the other even if one of them was very marginally improved. For example, if Lydia now found a 1 pence discount voucher for the Chinese takeaway, she would not suddenly come to prefer Chinese to pizza, despite the fact that she slightly prefers marginally-cheaper Chinese to marginally-more-expensive Chinese.²⁰

Now for the puzzle (due to Martin Peterson):

TAKEAWAY: Lydia has ordered the discounted Chinese but can switch her order to pizza. Assuming Lydia lacks a preference between these options, it seems permissible for her to switch. She does. Now, Lydia can switch again, to non-discounted Chinese (the discount has ended since her original order). Again, Lydia has no preference between these options, so it seems permissible to switch. Again, she does (Peterson 2007).

We have a puzzle. Lydia's acts seem permissible, as each involves trading between options that she lacks any preference between. Yet the sequence of acts seems problematic: while 1 pence isn't much, Lydia has given this up for no gain whatsoever. So, we have a mismatch between our evaluation of a sequence and our evaluation of the comprising acts. Further, as with Eve, Lydia's preferences are reasonable: it is reasonable to lack preferences between pizza and Chinese, with or without a tiny discount. So we have another potential analogue to our supererogation puzzles.

¹⁹ Formally, the previous case exhibited standard nontransitivity, while the current case exhibits transfinite nontransitivity (see Bartha et al. 2014). However, the important feature of SATAN'S APPLE is not this technical distinction but rather the fact that the story associated with the case makes the reasonableness of Eve's preferences clear.

²⁰ Lydia has, in a sense familiar from the decision-theoretic literature, incomplete preferences. For a discussion of such preferences, see Hare (2010).

3.3 The unproblematic-sequence solution

We can now look at how various people address these analogous puzzles. We start with Brian Hedden, who argues that the sequences in TAKEAWAY and SATAN'S APPLE are not *rationaly* problematic (Hedden 2015). Of course, it's unfortunate if Eve is banished and Lydia loses money. Nevertheless, says Hedden, there is nothing *normatively* problematic about the sequences of acts that lead to these outcomes. Given this, there is no conflict between our positive normative evaluation of the acts and our normative evaluation of the sequences.

Underpinning Hedden's claim is a time-slice view of rationality, on which an agent should treat different time slices of herself just as she treats distinct agents. Hedden notes that there are prominent cases (notably, the prisoners' dilemma) where a set of agents can all end up worse off if they all act rationally than if they all act irrationally. In other words, Hedden notes that distinct agents can end up poorly off as a result of their set of acts, without this revealing that the set is rationally problematic. On this time-slice view, the same is true of sequences of acts carried out by single agents: these sequences can be unfortunate without being rationally problematic.

The analogous solution to our supererogation puzzles is straightforward: it could be denied that the sequences in §2 are morally problematic. Thus we preserve HARMONY by denying, for example, that Scrooge's lifetime of omissions of supererogatory acts is morally problematic. Perhaps this sequence is unfortunate but, according to this solution, there's no *moral* issue with it. Call this the *unproblematic-sequence solution*.

Sadly, this solution is unconvincing. After all, Scrooge's behaviour does seem morally problematic, not merely unfortunate. We would need to be provided with strong grounds to deny that this was so. Further, Hedden's argument is far less forceful in the moral case than the rational. After all, while one arguably has no *rational* duty to coordinate with others in cases like the prisoner's dilemma, we often have *moral* duties to coordinate. Consequently, even if Scrooge should treat his different time slices in just the way he should treat other people, he is plausibly still morally required to coordinate with these other time slices. Therefore, a sequence of acts in which Scrooge fails to coordinate his actions is plausibly morally problematic, even given a time-slice view. So, the unproblematic-sequence solution is not compelling.

4 Resolving the puzzles: the sequence-sensitive solution

In the decision theory literature, various responses to the puzzles of SATAN'S APPLE and TAKEAWAY have been presented that, unlike Hedden, maintain the problematic nature of the sequences.²¹ These have a common thread: they hold that the place an act occupies in a sequence is of central normative relevance. It is this idea that we are interested in. We call the solution based on this idea the *sequence-sensitive solution*.

²¹ See, for example, Hare (2013), McClennen (1990), Wlodek Rabinowicz (1995), Rabinowicz (2000), Machina Mark (1989), Maher (1992) and Seidenfeld (1994).

4.1 Two distinctions

Let's start with two distinctions. First, our normative accounts can be made sensitive to sequences either by: (a) integrating sequential considerations into the individuation of acts; or (b) integrating sequential considerations into normative evaluations. Approach (a) declares that two acts are different to one another if they are embedded in different sequences or occur at different points in the same sequence. Given this, in paying attention to individual acts we automatically pay attention to the sequences in which they are embedded (because the sequence is part of what characterises the act). Approach (b) adopts a coarser-grained notion of acts: two acts can be tokens of the same type, despite being embedded in different sequences. However, when we evaluate acts, we evaluate them not only based on their own characteristics but also based on the sequences they are embedded in.²²

In practice, the two approaches are largely equivalent, as claims stated in the language of one can typically be restated in the language of the other. In this paper, we will therefore talk in terms of approach (b), as this approach makes particularly clear the importance of sequential considerations in normative reasoning.²³

So, we will argue that an act's place within a sequence is of central normative relevance. But what form does this relevance take? This brings us to our second distinction. There are two different approaches to answering this question. On the Act Approach, sequential considerations play a role in our normative evaluations of acts. On the Sequential Approach, on the other hand, it is not that sequential considerations impact our evaluations of acts, but that the full normative picture requires evaluating not just acts but sequences too. Let's consider these two approaches in detail.

4.2 The act approach

According to the Act Approach, which has previously been used to address puzzles like TAKEAWAY, act evaluations depends on sequential considerations.²⁴ This approach opens up the possibility that an act can be problematic in virtue of being part of a problematic sequence (poetically: the act can inherit the sins of the sequence). In turn, this makes room for a potential solution to our puzzling cases: we could argue that the morally-problematic nature of the sequences in these cases entails that at least one act in each sequence is itself problematic. This preserves HARMONY and so resolves the puzzling nature of these cases.

There are two ways of spelling out this solution. First, according to an approach we call *smooth*, if a sequence is problematic then every act in that sequence is problematic,

²² As should become clear, the way that sequential considerations play a role in evaluation might differ from the portrayal here. However, this complexity is addressed in the detailed discussion to follow, so we set it aside for now.

²³ For those who take there to be a significant difference between these two approaches, we have confidence that, even if not *all* the claims stated in the language of one can be restated in the language of the other, the particular claims we make in this paper can be.

²⁴ Central discussions of this strategy in the decision-theoretic literature include McClennen (1990), Rabinowicz (2000) and Wlodek Rabinowicz (1995). Related issues are discussed in discussions of Professor Procrastinate (see Jackson and Pargetter 1986).

to some extent.²⁵ Second, according to an approach we call *lumpy*, only some acts in a problematic sequence must themselves be problematic. Let's consider each view.

Smooth, when applied to our cases, turns out to have implausible implications, the extreme end of which is that it entails the impossibility of supererogatory action. Consider Scrooge: given *smooth*, the problematic nature of Scrooge's sequence entails that all his acts are morally problematic. Even if we take Scrooge's sequence to be morally bad (rather than *impermissible*), this means that each of Scrooge's omissions of supererogatory action is bad. And the claim that *any* omission of a supererogatory act is morally bad—let alone that *all* of Scrooge's are—is a controversial one, as many theorists insist that omissions of supererogatory actions are 'not bad', 'not wrong', 'morally neutral' and so on.²⁶ Further, consider the Saint. Given *smooth*, the fact that the Saint's sequence is morally bad entails that all her acts are (at least a little bit) bad. However, supererogatory acts are by definition *good* so it follows that the Saint has *never* performed a supererogatory act. This is implausible. *Smooth* fails.²⁷

Let's turn to *lumpy*, on which a sequence being problematic entails only that some of the constituent acts are morally problematic (and not that they all are). Well *lumpy* looks promising as a solution to THE COLLEAGUE. Here, it is fairly natural to think that some of Sara's decisions are supererogatory acts of generosity and some are problematic. Most straightforwardly, it might be thought that Sara's failure to give the final present is the act that calls out for censure. Why? Perhaps because it violates a duty of impartiality. Alternatively, perhaps it violates a duty to avoid meanness or cruelty: if we give gifts to 99 colleagues, it seems cruel to refuse to give a gift to the final colleague.²⁸ Either way, *lumpy* provided a promising solution to THE COLLEAGUE.

Unfortunately, when we turn to our other puzzling cases, *lumpy* begins to look problematically arbitrary. Given the implausibility of all the Saint's actions being morally bad, *lumpy* leads to the view that some of the Saint's acts are morally good but that other seemingly-equivalent acts are morally problematic. Similarly, *lumpy* entails that some of Scrooge's refusals to help others are morally problematic while others are not. Such judgements might appear hard to justify.

Still, *lumpy* can be defended. As a first step towards doing so, we note again the point raised in §2.2: sometimes the moral characteristics of a decision can depend on what decisions the agent has made in the past. Earlier, we pointed to promises, but

²⁵ *Smooth* may need to be restricted but, for simplicity, we focus on cases where each act in a sequence is of the same sort. Here, an unrestricted version of *smooth* suffices.

²⁶ Respectively: Mellema claims that acts whose omissions are morally bad acts are not supererogatory but, at best, quasi-supererogatory (1987); Heyd holds that only acts whose omissions are "not wrong" can be supererogatory (1982, p. 115); and Chisholm and Sosa, for example, claim that omissions of supererogatory acts are morally neutral (1966). Those who, like Driver, accept that it can sometimes be morally bad to refrain from supererogating only maintain that this holds in a special sort of morally-charged situation (1992, p. 288) and it would be extremely radical to claim that *every* occasion where Scrooge could perform a supererogatory act, he is in fact in a morally charged situation. Note, additionally, that the situation is not improved by moving to understand 'morally problematic' to indicated the stronger judgement of impermissibility, as this would entail that each one of Scrooge's omissions is impermissible and that is deeply implausible.

²⁷ The consideration of these implications of *smooth* and its subsequent failure puts pressure on Driver's smooth theory of the suberogatory (where suberogatory acts are those that are bad but permitted).

²⁸ Thanks to a reviewer for suggesting that a duty of this sort might be relevant here.

the phenomenon is far more general. For example, the first time I'm late to meet a friend is less problematic than the twentieth time I'm late. Likewise, perhaps Scrooge's refusal to donate is worse the twentieth time than the first, precisely because it occurs after nineteen refusals. Still, as we noted in our earlier discussion, it doesn't suffice to simply note this possibility. Instead, if we are to defend *lumpy* in this manner, we need to be given some grounds to think it's true that Scrooge's earlier acts influence the permissibility of the later acts.

In order to show this, we'll consider one way that the *lumpy* could be spelled out.²⁹ This version of *lumpy* appeals centrally to imperfect duties: duties that we must carry out at some point, but where we have some freedom regarding when or how we do so. As to how these duties play out, let's start simple. Imagine that an agent has some imperfect duty and knows they will only get to make ten decisions in their life. If the agent doesn't carry out the duty on any of the first nine occasions, then they know they will violate the duty if they don't carry it out on the 10th occasion. Under such circumstances (that is, given these previous decisions), it is natural to think that the agent is required to carry out the duty on the 10th occasion. So, here, the presence of an imperfect duty leads earlier decisions to influence the moral status of later decisions.

Now consider a more complex case. Once again, an agent has an imperfect duty, but this time, the agent is uncertain how many decisions they will get to make. In this case, there is no time at which the agent is sure that she would violate an imperfect duty if she made some decision. Nevertheless, each time the agent refuses to satisfy the imperfect duty, she comes to think it's a little more likely that she will die before satisfying the duty. As the agent comes to think this more and more likely, we might think some threshold is eventually passed. At this point, the risk of failing to satisfy the duty is high enough that the agent is required to satisfy it on the next occasion. Again, the presence of the imperfect duty makes the moral status of later decisions dependent on the decisions the agent previously made.

This account can make sense of the two cases that seemed to pose a problem for *lumpy* (SCROOGE and THE SAINT). Perhaps Scrooge has an imperfect duty of beneficence, and as a result, continual refusals to give eventually make Scrooge required to give. Similarly, perhaps the Saint has an imperfect duty of self-love, and as a result, continual refusals to act for herself eventually leave her required to do so. So this form of *lumpy* avoids a problematic arbitrariness in these cases.

Further, *lumpy* avoids the problem that arose for *smooth* and is compatible with the view that omissions of supererogatory acts are never morally bad. After all, *lumpy* divides Scrooge's choices into two categories. First, some of his choices will be morally unproblematic (and will plausibly involve omissions of supererogatory actions). Second, some of his choices will be morally problematic and we can simply deny that these omissions are of a supererogatory act. In neither case do we end up with the result that an omission of a supererogatory act is morally problematic, let alone that every such omission is. So *lumpy* avoids the problem that proved fatal for *smooth*. Overall, *Lumpy* is the most promising form of the Act Approach.

²⁹ Thanks to a reviewer for suggesting elements of the following account.

4.3 The sequential approach

Another version of the sequence-sensitive view, the Sequential Approach, is also familiar from the rational choice literature. According to this solution, we should reject HARMONY, accepting that a sequence of acts can be problematic without any constitutive act being problematic. On the Sequential Approach, then, we have local norms (applying to acts) and global norms (applying to sequences of acts), such that a sequence can be globally problematic, even if none of the acts that compose it are locally problematic. So the Sequential Approach takes it that reflecting on normative matters requires reflecting on sequences *as objects of evaluation*.

Given this, no contradiction arises from claiming that Eve's sequence of accepting every apple slice offered is rationally problematic (as it violates global norms) while her acts are all rationally permissible (as they satisfy local norms).³⁰ Likewise, there is no contradiction between claiming that Lydia's sequence is problematic in TAKEAWAY and claiming that her acts are all permissible (Hare 2013, Chapter 14).³¹ So, neither case raises a genuine puzzle.

The analogous solution in the context of supererogation is to accept global *moral* norms, such that a sequence can violate these while the constitutive acts satisfy local norms.³² On this view, HARMONY is false and so, we can accept that Scrooge's sequence of acts is problematic, while each act is unproblematic (and can make sense of the other cases in a similar manner).

This solution comes in two flavours, distinguished by how they account for supererogation.³³ On the first, whether an act is supererogatory is a local matter. So, each of the Saint's acts (say) are supererogatory, because they are locally permissible and locally better than a permissible alternative. Further, this claim holds even though the sequence of acts is problematic. This account allows us to retain traditional views of when an act is supererogatory, but this comes at a cost: it undermines the significance of the supererogatory. After all, on this account the fact that an act is supererogatory doesn't guarantee that it is part of a globally-acceptable sequence. So, choosing a supererogatory act may lead an agent to violate global norms. Consequently, there will be cases where an agent will, if moral, avoid a supererogatory act. A cost indeed. Yet we doubt that this cost is substantial enough to justify abandoning the view. After all, the existence of our puzzling case reveals that some change to our thinking about supererogation is needed.

Still, there's an alternative: perhaps acts are supererogatory only if they both satisfy local norms *and* play an appropriate global role. This departure from the traditional,

³⁰ Arntzenius, Elga, and Hawthorne adopt something close to this view, by rejecting the Deal Agglomeration Principle (2004).

³¹ He adopts an account of this sort in response to cases like TAKEAWAY. Hare treats sequences as composite acts and denies that there is a simple relationship between the status of a composite act and the status of the acts it is composed from.

³² In the sort of cases we're discussing, there will always be some sequence that is globally permissible and is made up of locally-permissible acts (briefly: due to the optionality of the supererogatory). So, the view here does not lead to moral dilemmas. Further, even if global and local norms clashed, dilemmas need not follow: perhaps the norms can be balanced to determine what is moral all-things-considered or perhaps, for any circumstance, one set of norms trumps the other (see Lazar and Lee-Stronach).

³³ Indeed, the two versions could be combined: perhaps supererogation can arise at multiple levels.

local view also comes at a cost: it will often make it difficult to determine whether an action is supererogatory, as doing so requires looking beyond the current choice to past and potential future choices.³⁴ Again, this is a genuine cost. Again, it is a cost that can be borne: sometimes moral evaluations simply are hard to carry out and so it is hardly fatal for an account that it entails such difficulties.

In either case, a problem arises for the sequential solution from some plausible claims about conditional obligations, where these are the obligations we would have if certain conditions were met.³⁵ For concreteness, we'll focus on the case of Scrooge, and we'll assume that Scrooge knows he will face exactly ten opportunities to give. Here's a first claim about conditional obligations: Scrooge has a conditional obligation to give on the tenth occasion if he hasn't given on the first nine occasions. Here's a second claim about conditional obligations: if an agent has an obligation to O conditional on C then the agent also has an unconditional obligation to O if the conditions specified in C are met. It follows from these two plausible claims that Scrooge has an obligation to give on the tenth occasion if he hasn't given on the first nine. This contradicts the claim, endorsed by the sequential approach, that none of Scrooge's acts are problematic in and of themselves.

A natural response on behalf of the sequential view is simply to deny that Scrooge has the specified conditional obligation. After all, why think he does? Plausibly, we think this because we accept that Scrooge has an imperfect duty to give and think that imperfect duties provide us with conditional obligations. Yet the proponent of the sequential view can simply deny this. On the sequential view, imperfect duties can be thought of as operating on the global level and not on the local level at all. Imperfect duties, then, provide obligations to avoid certain sequences of actions but do not provide obligations (conditional or otherwise) to carry out particular actions. So, just as with the act approach, the sequential approach provides a promising way to spell out the sequence-sensitive solution.

5 Where we find ourselves

We endorse the sequence-sensitive solution; we think you should too. After all, both forms of this solution resolve our puzzles. Further, they do so in a natural way. It was precisely the sequential characteristics of these puzzles that were problematic. It's natural to account for this by allowing sequential considerations a role in moral evaluations.

As to which form of the sequence-sensitive solution should be adopted, this is a harder matter to settle. Indeed, given that both the Act Approach and the Sequential Approach resolve our puzzles, we doubt that these puzzles alone support one form over the other. Instead, which should be adopted plausibly depends on broader questions about the moral world: questions that are too broad to answer here. For example, do we have independent reasons to accept the existence of global norms (or, indeed, to doubt their plausibility)? Likewise, do we have independent reasons to accept that

³⁴ A, perhaps more plausible, alternative: to the agent's *current beliefs* about her past and future choices.

³⁵ We owe this objection to a reviewer for the journal.

local norms sometimes depend on sequential considerations (or, indeed, reasons to doubt this)?

The point is that one should not adopt either form of the sequence-sensitive solution independently of their general moral views. Instead, one should adopt the solution that is in closest conformity with these broader views. Such a process will lead different readers to different solutions. For example, a reader who thinks Scrooge clearly has a conditional obligation to give on the tenth occasion will be pushed towards the act approach; someone who denies this might find more sympathy with the sequential approach. Different views about the broader moral world support different conclusions about the narrow issue under discussion here.

Still, for what it's worth, we adopt the Sequential Approach. Our core reasoning is simple: we think it natural to morally evaluate sequences. Indeed, we think it strange to deny that doing so is possible. Given this view of the moral world, the Sequential Approach naturally resolves our puzzle cases. After all, Scrooge's sequence does seem morally problematic and his individual acts unproblematic, so an account that accepts this is promising.

Further, the account can also make sense of cases where it seems even clearer that the sequence is problematic and yet, contrary to the Act Approach, no component act is. Consider:

POCKET MONEY: Kwame is deciding how much weekly pocket money to give to his twin daughters. Any amount from nothing to £10 is an acceptable amount to give as pocket money.³⁶ Given that any amount is acceptable, Kwame decides to give one daughter £1 and the other £10.

Here Kwame's sequence (giving one daughter £1 and the other £10) is problematic. The problem with the sequence can be equally well solved by levelling up as by levelling down: that is, by giving the one daughter more or the other less, so that they receive the same amount. Therefore, it seems wrong to single out the latter act as wrong; and equally wrong to single out the former. Unlike in the COLLEAGUE case, where it is the refusal to give the final gift that seems problematic, in cases like POCKET MONEY neither individual action is obviously the source of the issue. Of course, it might be argued, *à la smooth*, that both acts are equally wrong. However, the Sequential Approach is able to make sense of those of us with the intuition that it is mistaken to attribute blame to every act whenever no particular act can be singled out.

This core reasoning is then bolstered by reflection on a distinct puzzle from the supererogation literature, which the Sequential Approach provides a particularly satisfying solution to. The puzzle arises because some apparently supererogatory actions do not seem to call for praise. This contradicts a longstanding view that supererogatory actions are, without exception, praiseworthy.

It is a matter of dispute which actions have this puzzling form, but one plausible candidate is proposed by Jason Kawall:

³⁶ We take this as a stipulated assumption, setting aside worries that no pocket money for children is somehow a problem but also that £10 is somehow excessive. The amounts can easily be changed, it is the structure of the example that is important.

What are we to make of the selfless actions of a woman whose self-esteem has been crippled by a verbally abusive husband and a traditional society which teaches that women are first and foremost caregivers? Or consider a cult member who does not lack self-esteem but who has placed the cult leader on a high pedestal such that he would do anything (thus going beyond duty) to benefit the leader (Kawall 2003, p. 495).³⁷

Here there is an inconsistency between three plausible claims: (1) the selfless acts of the cult member (say) are supererogatory; (2) supererogatory acts are praiseworthy; (3) the way that the cult member acts is not praiseworthy.

Unsurprisingly, traditional responses to such cases deny at least one of (1)–(3). For example, Nancy Stanlick rejects (2), claiming that “supererogatory acts are, in fact, not necessarily praiseworthy [or] morally commendable” (Stanlick 1999, p. 210). An alternative response is the one Kawall proposes: to reject (1) and so deny that the cult member’s selfless acts are supererogatory. In either case, the solution comes at a cost, as each of (1)–(3) is plausible.

The Sequential Approach can avoid this cost; it can allow us to retain (1)–(3), by distinguishing the levels at which each applies. (1), it could be said, applies at the local level: considered locally, the cult member acts in a supererogatory manner on each occasion.³⁸ (2) then holds, again at this local level: it entails that each of the cult member’s supererogatory acts are locally praiseworthy. Finally, this does not contradict (3), which applies at the global level: the cult member’s sequence of acts is not globally praiseworthy. Indeed, at a global level it calls for criticism. So we can make sense of Stanlick’s intuition that supererogatory acts are sometimes morally criticisable, without abandoning (2), by revealing that some acts can be locally supererogatory and thus locally praiseworthy, while being morally criticisable on a global level. So the Sequential Approach can make sense of cases where actions seem to be supererogatory but not praiseworthy without rejecting any of the three plausible claims.

6 Conclusions

Myopic accounts of supererogation focus only on an agent in the moment of decision and ignore the acts that the agent has performed in the past or expects to perform in the future. However, the plausibility of such accounts as *complete* accounts of the supererogatory comes into question when we consider various puzzling cases. In response to this, we have drawn on an analogy between these puzzles and well-known

³⁷ For a related discussion, see also Hampton (1993, p. 1). Now, some readers might have doubts about the particular cases that Kawall describes (is every individual action of the cult member and deferential wife really acceptable?). However, what is important is not the details of the example but that our account can make sense of cases like these commonly thought to be problematic in the supererogation literature where agents seem engaged in behaviour that isn’t praiseworthy and yet each individual act seems to be. By differentiating local and global norms, our account can make sense of these cases without falling into the dilemma previously thought to be inevitable: that either these acts are not in fact supererogatory or we supererogatory acts are not in fact praiseworthy.

³⁸ This commits us to the view that supererogation arises at the local level. Still, this need not require denying that supererogation can also arise at the global level.

puzzles from the rational-choice literature to find a new account of supererogation. This has led us to the sequence-sensitive solution, according to which our accounts of supererogation must pay attention not simply to actions considered in isolation but also to the sequences comprised of these actions. We endorse this view.

We also endorse a particular version of the sequence-sensitive solution: the Sequential Approach. However, our endorsement here is weaker: this approach fits with our own broader views of the moral world. Yet we acknowledge that others, with different views of broader matters, might instead be led to the Act Approach.

In either case, the important point is this: the ethical world cannot be understood via a myopic focus on just the current decisions that we face. Rather, the ethical world is inextricably entwined with temporally-extended considerations. In order to live, and understand, our moral lives, we must look beyond our individual decisions and consider the patterns in which they are embedded.

Acknowledgements We thank all the participants of the *Ethics and Risk* conference at ANU 2018, as well as the anonymous reviewers of this journal for their invaluable feedback.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Armtenius, F., Elga, A., & Hawthorne, J. (2004). Bayesianism, infinite decisions, and binding. *Mind*, 113(450), 251–283.
- Bartha, P., Barker, J., & Hajek, A. (2014). Satan, Saint Peter and Saint Petersburg: Decision theory and discontinuity at infinity. *Synthese*, 191(4), 629–660.
- Benn, C. (2017). Supererogation, optionality and cost. *Philosophical Studies*, 175(10), 2399–2417. <https://doi.org/10.1007/s11098-017-0965-7>.
- Brandt, R. (1979). *A theory of the good and the right*. Oxford: Clarendon Press.
- Chisholm, R. M., & Sosa, E. (1966). Intrinsic preferability and the problem of supererogation. *Synthese*, 16, 321–331.
- Davidson, D., McKinsey, J. C. C., & Suppes, P. (1955). Outlines of a formal theory of value, I. *Philosophy of Science*, 22(2), 140–160.
- Driver, J. (1992). The subrogatory. *Australasian Journal of Philosophy*, 70(3), 286–295.
- Hajek, A. (2009). Dutch book arguments. In P. Anand, P. K. Pattanaik, & C. Puppe (Eds.), *The handbook of rational and social choice: An overview of new foundations and applications* (pp. 173–195). Oxford: Oxford University Press.
- Hampton, J. (1993). Selflessness and the loss of self. *Social Philosophy and Policy*, 10(1), 135–165.
- Hare, C. (2010). Take the sugar. *Analysis*, 70(2), 237–247.
- Hare, C. (2013). *The limits of kindness*. Oxford: Oxford University Press.
- Hedden, B. (2015). *Reasons without persons*. Oxford: Oxford University Press.
- Heyd, D. (1978). Ethical universalism, justice, and favouritism. *Australasian Journal of Philosophy*, 56(1), 25–31.
- Heyd, D. (1982). *Supererogation: Its status in ethical theory*. Cambridge studies in philosophy Cambridge: Cambridge University Press.

- Horgan, T., & Timmons, M. (2010). Untying a knot from the inside out: Reflections on the ‘paradox’ of supererogation. *Social Philosophy and Policy Foundation*, 27(02), 29–63. <https://doi.org/10.1017/S026505250999015X>.
- Jackson, F., & Pargetter, R. (1986). Oughts, options, and actualism. *Philosophical Review*, 95(2), 233–255.
- Kawall, J. (2003). Self-regarding supererogatory actions. *Journal of Social Philosophy*, 34(3), 487–498. <https://doi.org/10.1111/1467-9833.00196>.
- Lazar, S., & Lee-Stronach, C. (n.d.). Axiological absolutism and risk. *Nous* (forthcoming).
- Lowe, S. C. (2009). Ebenezer scrooge—Man of principle. *Think*, 8(23), 27–34. <https://doi.org/10.1017/S1477175609990030>.
- Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(December), 1622–1668.
- Maher, P. (1992). Diachronic rationality. *Philosophy of Science*, 59(1), 120–141.
- Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25, 205–207.
- McClennen, E. F. (1990). *Rationality and dynamic choice: Foundational explorations*. Cambridge: Cambridge University Press.
- McClennen, E. F. (1997). Pragmatic rationality and rules. *Philosophy & Public Affairs*, 26(3), 210–258.
- Mellema, G. (1987). Quasi-supererogation. *Philosophical Studies*, 52, 141–150.
- Peterson, M. (2007). Parity, clumpiness and rational choice. *Utilitas*, 19(4), 505–513.
- Portmore, D. W. (2013). Perform your best option. *Journal of Philosophy*, 110(8), 436–459.
- Portmore, D. W. (2016). Transitivity, moral latitude and supererogation. *Utilitas*. <https://doi.org/10.1017/S0953820816000364>.
- Quinn, W. S. (1990). The puzzle of the self-torturer. *Philosophical Studies*, 59(1), 79–90.
- Rabinowicz, W. (1995). To have one’s cake and eat it, too: Sequential choice and expected-utility violations. *Journal of Philosophy*, 92(11), 586–620.
- Rabinowicz, W. (2000). Money pump with foresight. In M. J. Almeida (Ed.), *Imperfectible harms and benefits* (pp. 123–154). Berlin: Kluwer Academic Publisher.
- Rawls, J. (1999). *A theory of justice* (Rev ed.). Cambridge, MA: The Belknap Press of Harvard University Press.
- Raz, J. (1975). Permissions and supererogation. *American Philosophical Quarterly*, 12(2), 161–168.
- Seidenfeld, T. (1994). When normal and extensive form decisions differ. In D. Prawitz, B. Skyrms, & D. Westerståhl (Eds.), *Logic, methodology and philosophy of science* (pp. 451–463). Amsterdam: Elsevier.
- Stanlick, N. A. (1999). The nature and value of supererogatory actions. *Journal of Social Philosophy*, 30(1), 209–222. <https://doi.org/10.1111/0047-2786.t01-1-00013>.
- Wolf, S. (1982). Moral saints. *The Journal of Philosophy*, 79(8), 419–439.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.