

Aggregation and Reductio¹

Patrick Wu

Joe Horton argues that partial aggregation yields unacceptable verdicts in cases with risk and multiple decisions. I begin by showing that Horton's challenge does not depend on risk, since exactly similar arguments apply to riskless cases. The underlying conflict Horton exposes is between partial aggregation and certain principles of diachronic choice.

I then provide two arguments against these diachronic principles: (i) they conflict with intuitions about parity, prerogatives, and cyclical preferences, and (ii) they rely on an odd assumption about diachronic choice. Finally, I offer an explanation, on behalf of partial aggregation, for why these diachronic principles fail.

Consider two questions:

1. You can either save a million people from becoming quadriplegics or a single person from death. Should you save the one?
2. You can either save a zillion people from headaches or a single person from death. Should you save the one?

Partial aggregationists answer no to the first question and yes to the second. On this view, when harms are relevantly similar (e.g. becoming quadriplegic and death), we can aggregate, but when harms are very different (e.g. headaches and death) we cannot. Full aggregationists, by contrast, answer no to both questions. They aggregate even when one person faces a far greater harm. If a zillion is large enough, they conclude that you ought to prevent the headaches.

¹ For discussion, I thank Zach Barnett, Daniel Muñoz, and Aidan Penn. For feedback on drafts, I thank Tez Clark, Bar Luzon, Stephan Pohl, Samuel Scheffler, and two associate editors of *Ethics*. Special thanks to Cian Dorr and Daniel Viehoff.

Joe Horton has recently provided three excellent arguments against partial aggregation (PA). This paper, in response, has three aims: (i) to identify the principles in Horton's arguments that PA must reject, (ii) to argue against those principles, and (iii) to explain why such principles do not apply to PA.

The paper proceeds as follows. I present Horton's three arguments, showing that each relies on a bespoke diachronic principle (§1). Next, I show that each diachronic principle alone reduces PA to absurdity (§2). So PA must reject all three diachronic principles. Finally, I give general reasons for doubting these principles (§3), and I explain why they fail for PA in particular (§4).

Section 1: Horton's Arguments

Consider three cases:²

Case 1

Someone has a headache. You can cure the headache by snapping your fingers, but if you do so, there is a one in a zillion chance that some other person will die. Should you snap?

PA Intuition 1: You are required to snap. People routinely expose some to tiny chances of death in the course of helping others, e.g. when driving to get someone aspirin. Sometimes, such behavior is required.

Case 2

² These cases and intuitions come from Horton 2020; I have modified the cases slightly to remove the villain, a potential confounder for our intuitions.

To your left, there are a zillion people with headaches. To your right, there are a zillion others. You can cure the headaches by snapping your fingers, but if you do so, then one random member of the zillion to your right will die. Should you snap?

PA Intuition 2: you are not permitted to snap. Headaches, when compared to lives, are insufficiently weighty; saving one life would outweigh curing even a zillion headaches.

Case 3

You find yourself trapped in a long corridor; you see a zillion rooms to your left and a zillion to your right. Each room on the left contains one person with a headache, and each on the right contains another person. Between each pair of rooms, there is a big button. If you press the button, then the headache to your left will be cured and the person to your right will receive a ticket for a lottery with a zillion tickets. After you pass by all the rooms (as you must), a random ticket will be drawn and the winner (if any) killed.

Should you press the first button? Should you press the second button? ... Should you press the zillionth button?

In Case 3, unlike with the others, there are no direct intuitions about what the agent ought to do. There are, after all, zillions of possible sequences which all differ in small ways. But as Horton notes, pressing the first button in Case 3 seems "exactly like" snapping in Case 1.³ Since there you may snap, here you may press. Similar remarks apply to the other buttons. Thus, for each button, you are permitted to press it. In short:

Moral Equivalence 1 (ME 1)

³ Horton 2020, pg. 517.

In virtue of some similarities S , if you are permitted in Case 1 to (not) snap, then you are permitted, in Case 3, to (not) press the n th button.

I follow Horton in leaving S unspecified. Instead, I will use pairs of cases that seem to be morally equivalent for the same reasons as apply to Horton's, so that any plausible S will either apply to both pairs of cases, or to neither pair of cases. Nonetheless, it is worth identifying some features of Case 1 and Case 3 that make their moral equivalence seem plausible, in order to verify that those features are shared by the cases I will present later. Here are three such features:

1. Snapping (not snapping) in Case 1 has the same non-normative effect as pressing the n th button (not pressing the n th button) in Case 3.
2. Regardless of one's previous choices in Case 1 (Case 3), one faces the choice of whether to snap (press the n th button) and the non-normative effects of snapping (pressing the n th button) are unchanged by one's previous choices.
3. Regardless of one's subsequent choices in Case 1 (Case 3), the non-normative effects of one's past choice to snap or not (press the n th button or not) will remain unchanged.

Feature (1) is clearly required for moral equivalence. Features (2) and (3), while vacuously satisfied for Case 1, help make each choice in Case 3 seem "independent" of the others. As a result, it is plausible that one's choices in Case 3 should match one's choices in the one-off variant exemplified by Case 1, just as ME 1 asserts. With ME 1 and PA Intuition 1, we get:

Lemma

In Case 3, for each button, you are permitted to press it.

It is important to be careful about what Lemma says. In general, one can distinguish questions about:

1. Individual Choices: What should an agent do at some particular choice-point?
2. Sequences: Which sequence of actions should an agent perform?⁴
3. Outcomes: Which sequence's final outcome is most choiceworthy?

Lemma concerns the permissibility of individual choices. It says, for example, that it is permissible to press button 1. And it says that at the choice-point after one has pressed button 1, not pressed button 2, ..., and pressed button 423, it is permissible to press button 424. But Lemma does not say anything about sequential permissibility (e.g. that it is permissible to press button 1 + press button 2 + ... + press button zillion) or outcome permissibility (e.g. that it is permissible to choose the outcome of pressing all the buttons over the outcome of pressing none of the buttons). To draw such implications from Lemma, one needs further principles.

Horton next argues that PA Intuition 2 and Lemma conflict. This is because pressing every button in Case 3 is "in all morally relevant respects equivalent" to snapping in Case 2.⁵ Such thoughts are plausible, but it helps to make the reasoning explicit.

There seem to be two thoughts at play. The first is that, if each action in some sequence is individually permissible, then surely the sequence of actions is permissible.

More formally:

⁴ You may think that this question does not make sense as a separate question, distinct from (1) and (3). Perhaps you think that completing a permissible series of actions is just a matter of performing a series of actions such that each member in the series was permissibly chosen at the appropriate junction. This view reduces (2) to (1). Or you may think that completing a permissible series of actions is just a matter of ending up with a permissible outcome. This view reduces (2) to (3). Since both views are plausible, I remain neutral about how best to understand (2).

⁵ Horton 2020, pg. 517.

Individual Choices to Sequences I (ICS I)

Suppose you know that you will face n pairwise choices between X_i and Y_i . If, for each i , you are permitted to choose X_i over Y_i , then you are permitted to choose the sequence $X_1 + \dots + X_n$.

Thus, if it is individually permissible to press each button, then the sequence of (press button 1 + press button 2 ... + press button zillion) is permissible.

The second thought is that, if some sequence is permissible to choose even though another sequence is available, then surely the outcome of the first sequence is permissible to choose over the outcome of the second sequence.⁶ For if not, why would the first sequence be permissible? In other words:

Sequences to Outcomes I (SO I)

Suppose you know that you will face n pairwise choices between X_i and Y_i . If you are permitted to choose the sequence $X_1 + \dots + X_n$, then in a separate choice between the outcome of $X_1 + \dots + X_n$ and the outcome of $Y_1 + \dots + Y_n$, you are permitted to choose the former.

Thus, if it is permissible to choose the sequence (press button 1 + ... + press button zillion), then it is permissible to choose the outcome of (press button 1 + ... + press button zillion) over the outcome of (don't press button 1 + ... + don't press button zillion). Now, Horton's first argument:

Argument I

- I. In Case I, one may snap. (PA Intuition I)

⁶ One might, instead, think that outcomes should constrain sequential permissibility: if some outcome is worse than another, then any sequence that leads to the former outcome is not permissible when some sequence that leads to the latter is available. But that is just the contraposited version of this thought.

2. In Case 2, one must not snap. (PA Intuition 2)
3. If in Case 1, one may snap, then in Case 3, for each button, one may press it. (ME 1)
4. If in Case 3, for each button, one may press it, then in Case 3, one may choose the sequence (press button 1 + ... + press button zillion). (ICS 1)
5. If in Case 3, one may choose the sequence (press button 1 + ... + press button zillion), then in Case 2, one may snap. (SO 1)
6. Thus, in Case 2, one may snap. (By 1, 3, 4, 5)
7. Contradiction. (By 2, 6)

Section 1.1: Horton's Second and Third Arguments

Horton's other arguments rely on subtly different principles. Argument 1 goes from permissions (in Case 1) to permissions (in Case 2). By contrast, Argument 2 goes from requirements to requirements, and Argument 3 goes from requirements to permissions. It may seem pedantic to care about these differences, but as §3 discusses, the principles involved in each argument have very different implications in other cases. So it is worth distinguishing the arguments.

To get the remaining arguments, consider these schemas:

Moral Equivalence (ME)

In virtue of some similarities S, if you are ___ in Case 1 to (not) snap, then you are ___, in Case 3, to (not) press the nth button.

Individual Choices to Sequences (ICS)

Suppose you know that you will face n pairwise choices between X_i and Y_i . If, for each i , you are ___ to choose X_i over Y_i , then you are ___ to choose the sequence $X_1 + \dots + X_n$.

Sequences to Outcomes (SO)

Suppose you know that you will face n pairwise choices between X_i and Y_i . If you are ___ to choose the sequence $X_1 + \dots + X_n$, then in a separate choice between the outcome of $X_1 + \dots + X_n$ and the outcome of $Y_1 + \dots + Y_n$, you are ___ to choose the former.

To get ME 2, ICS 2, and SO 2, substitute “required” for every blank. To get ME 3, ICS 3, and SO 3, substitute “required” for the first three blanks and substitute “permitted” for the rest. The resulting principles are plausible and can be motivated by speeches similar to the ones given above. With them, Horton can complete:

Argument 2

1. In Case 1, one must snap. (PA Intuition 1)
2. In Case 2, one may refrain from snapping. (PA Intuition 2)
3. If in Case 1, one must snap, then in Case 3, for each button, one must press it. (ME 2)
4. If in Case 3, for each button, one must press it, then in Case 3, one must choose the sequence (press button 1 + ... + press button zillion). (ICS 2)
5. If in Case 3, one must choose the sequence (press button 1 + ... + press button zillion), then in Case 2, one must snap. (SO 2)
6. Thus, in Case 2, one must snap. (By 1, 3, 4, 5)
7. Contradiction. (By 2, 6)

and

Argument 3

1. In Case 1, one must snap. (PA Intuition 1)
2. In Case 2, one must not snap. (PA Intuition 2)
3. If in Case 1, one must snap, then in Case 3, for each button, one must press it. (ME 3)
4. If in Case 3, for each button, one must press it, then in Case 3, one may choose the sequence (press button 1 + ... + press button zillion). (ICS 3)
5. If in Case 3, one may choose the sequence (press button 1 + ... + press button zillion), then in Case 2, one may snap. (SO 3)
6. Thus, in Case 2, one may snap. (By 1, 3, 4, 5)
7. Contradiction. (By 2, 6)

Section 2: PA's Reply to Horton's Arguments

How should PA respond to Horton's arguments? There seem to be two options: (i) reject these verdicts about risky cases, or (ii) reject some of the principles connecting individual, sequential, and outcome permissibility. But as I will show, Horton's arguments can be strengthened to avoid reliance on risky cases. More precisely, define for $i = 1, 2, 3$:

Diachronic Principle i (DP i)

Moral Equivalence $i \wedge$ Individual Choices to Sequences $i \wedge$ Sequences to Outcomes i

This section shows that each DP i alone reduces PA to absurdity. So PA can only take the second option.

Start by assuming that PA delivers the following representative verdicts:⁷

⁷These nice examples come from Horton 2018.

PA Verdicts

Ceteris paribus:

1. In choosing between 20n arms and n lives, you are required to save the arms.
2. In choosing between 20n fingers and n arms, you are required to save the fingers.
3. In choosing between a life and n fingers, you are required to save the life.

I assume these verdicts for concreteness; they are really stand-ins for whatever corresponding verdicts one's favored variant of PA delivers. Now, consider:

Optimal Beneficence

Let X and Y be disjoint populations. You can only save one of them. Each population is collectively at risk of losing:

- X: 2 lives, 80 arms, 3,200 fingers
- Y: 1 life, 40 arms, 1,600 fingers

You are required to save X.

Any plausible PA view will accept something like Optimal Beneficence. Finally, consider:

Islands

There are six islands (A, B, C, D, E, F) with disjoint populations. You have to make three choices about who to save. You know in advance that you will have to make all three choices.

- Monday: you choose between saving two lives (on island A) and saving 40 arms (on island D)
- Tuesday: you choose between saving 80 arms (on island B) and saving 1,600 fingers (on island E)

- Wednesday: you choose between saving 3,200 fingers (on island C) and saving 1 life (on island F)

Which island should you choose on Monday? On Tuesday? On Wednesday?

I start with DP 1. Suppose, for contradiction, that DP 1 was true. Then, by PA Verdicts, in the following one-off variant:

Case 4

You can either save two lives (on island A) or 40 arms (on island D).

What should you do?

You may save D. Recall that ME 1 is part of DP 1. Whatever similarities exist between snapping in Case 1 and pressing a button in Case 3 also seem present between saving D in Case 4 and saving D in Islands. For example, recall the features considered earlier: (i) the effects of the actions are the same, (ii) neither case permits earlier actions to alter the effects or availability of later actions, and (iii) neither case permits later actions to change the effects of earlier ones. All those features hold here too. So there is no reason, once one grants ME 1, for denying:

ME 1*

In virtue of some similarities S, if you are permitted in Case 4 to (not) save D, then you are permitted, on Monday in Islands, to (not) save D.

Thus, on Monday, you may save D. Exactly similar arguments apply to get a permission, on Tuesday, to save E and a permission, on Wednesday, to save F. But then, by ICS 1, you may choose the sequence (D+E+F) in Islands. Then, by SO 1, in a separate choice between the outcome of (D+E+F) and the outcome of (A+B+C), you may save the former. That is, in a

choice between saving 1 life + 40 arms + 1600 fingers and saving 2 lives + 80 arms + 3200 fingers, you may save the former. That contradicts Optimal Beneficence.

Exactly similar arguments apply to DP 2 and DP 3. Thus, PA must reject all three principles.

Section 3: Against the Diachronic Principles

Horton's arguments force a choice between PA and the DPs. This section presents two reasons for doubting the DPs. First, they conflict with other normative phenomena: DP 1 conflicts with parity and supererogation, DP 2 conflicts with personal prerogatives, and DP 3 conflicts with cyclical preferences. Second, the DPs rely on a common assumption which, when made explicit, appears dubious.

Those familiar with the thought that these normative phenomena conflict with diachronic consistency principles may prefer to skip §3.1 – §3.3. (I include those sections only because standard discussions of such conflicts do not appeal to the DPs.⁸)

Section 3.1: DP 1, Parity, and Supererogation

I show how DP 1 conflicts with parity and supererogation.⁹

What is parity? Sometimes, I have no strict preference between two goods. For example, if asked to choose between a cup of coffee (rich, nutty, pleasantly acidic) and tea (bright, vegetal, sweet), I do not prefer either to the other. Furthermore, my attitude is not

⁸ Some discussions are unclear about which diachronic consistency principles are at play (e.g. Chang 1997/2005). Others appeal to Peter Hammond's diachronic consistency conditions (Bader 2019), dominance principles (Dougherty 2014), or backwards induction principles (Gustafsson 2020).

⁹ Conflicts between parity and diachronic consistency principles are well-known (Chang 1997, pgs. 10 – 11; 2005, pgs. 346 – 347). Contemporary accounts of supererogation (e.g. Muñoz 2021; Parfit 2011, pgs. 137 – 139) are often structurally similar to parity, and as a result, the problem generalizes (Bader 2019).

indifference since I continue to vacillate even after “slight sweetenings.” For example, I do not prefer (coffee with \$1) to tea, nor do I prefer (tea with \$1) to coffee. In such cases, coffee and tea are “on a par.”¹⁰

Suppose that you share these preferences, and that you prefer, *ceteris paribus*, more money to less. Consider:

Case 5

You can either choose (coffee with \$1) or tea.

Given the assumptions, you are permitted to choose either option. Similarly, in:

Case 6

You can either choose (tea with \$1) or coffee.

You may choose either. But now consider:

Case 7

You know that you will have to make two choices.

Monday: choose between (coffee with \$1) and tea

Tuesday: choose between (tea with \$1) and coffee

What should you choose on Monday? On Tuesday?

I also assume that you have no preferences for when you consume coffee/tea. You do not, for example, prefer consuming coffee on Monday over consuming coffee on Tuesday.

Suppose for contradiction that DP 1 is true. Recall that ME 1 is part of DP 1. Again, whatever features Case 3 and Case 1 share, it seems Case 7 and Case 5 share them as well. So there seems to be no reason for now denying:

¹⁰ In general, X and Y are on a par iff [I do not strictly prefer X to Y, I do not strictly prefer Y to X, and there exists Z such that (I strictly prefer Z to X and I do not strictly prefer Z to Y)] or Y and X are on a par. See also Chang 2002 and Hare 2010.

ME I**

In virtue of some similarities S, if you are permitted in Case 5 to (not) choose tea, then you are permitted, on Monday in Case 7, to (not) choose tea.

Thus, on Monday, you may choose tea. Similarly, on Tuesday, you may choose coffee. So, by ICS I, you may choose the sequence (tea + coffee). So, by SO I, in a separate choice between the outcome of (tea + coffee) and the outcome of (coffee with \$I + tea with \$I), you may choose the former. But that is not true since, by hypothesis, you prefer more money to less. Contradiction.

DP I also conflicts with supererogation. Consider:

No Sacrifice Needed

Suppose you must choose between: (losing a zillionth of a dollar + some stranger gains an extra dollar) and (keeping all of your resources + stranger gains nothing).

Then, you are permitted to keep all of your resources.

Intuitively, one is not required to do anything (not even give up a zillionth of a dollar) to help a stranger get just one dollar. However, the following is also plausible:

Sacrifice Needed

Suppose you must choose between: (losing a dollar + a zillion strangers gain an extra dollar each) and (keeping all of your resources + strangers gain nothing). Then, you are required to give the dollar.

When losing a dollar can benefit so many others to just the same degree, one is required to give.

If one accepts both judgments, then DP I again leads to contradiction. Imagine facing a sequence of a zillion cases as described in No Sacrifice Needed. By ME I and No

Sacrifice Needed, for each choice, you are permitted to keep your zillionth. Then, by ICS 1, you are permitted to choose (keep the zillionth + ... + keep the zillionth). Then, by SO 1, you are permitted to choose the outcome of that sequence (i.e. you keep your resources and the strangers get nothing) over the outcome of the sequence where you give your zillionths away (i.e. you lose a dollar and the zillion strangers each get a dollar). This contradicts Sacrifice Needed.

Section 3.2: DP 2 and Prerogatives

I show how DP 2 conflicts with personal prerogatives. Consider:

Small Sacrifice

Suppose you will suffer n seconds of great pain. By suffering for an extra second, you can spare a hundred people an hour of great pain. For any natural number n , you must choose to suffer the extra second.

In any such one-off choice, it seems one must suffer the extra second. But now consider:

Big Sacrifice

By suffering great pain for many decades, you can spare n people each an hour of great pain. For any choice of n , you are permitted to refuse the decades of suffering.

Intuitively, there are limits to what morality demands of people; it cannot, for example, require one to throw away much of one's life just to spare others an hour of pain.

If one accepts both judgments, then DP 2 leads to contradiction. Imagine facing a sequence of 20,000 choices as described in Small Sacrifice. By ME 2 and Small Sacrifice, you are required to accept each second of suffering. Then, by ICS 2, you are required to choose the sequence (suffer the extra second + ... + suffer the extra second). Then, by SO 2, you are

required to choose the outcome of (suffer the extra second + ... + suffer the extra second) over the outcome of (decline the extra second + ... + decline the extra second). So you are required to accept 20,000 seconds (~ 5.5 decades) of suffering in order to spare 20,000 strangers each an hour of suffering. This contradicts Big Sacrifice.

Section 3.3: DP 3 and Cyclical Preferences

I show how DP 3 conflicts with cyclical preferences. Consider:

Extra Turn¹¹

Suppose you are connected to a machine that provides electrical shocks. Its values range from 0 to 1,000 where 0 yields no shocks and 1,000 yields excruciating shocks. For any n , there is barely any difference in pain between n and $n+1$.

Right now, the machine is at p . If you press the button, then the machine goes to $p + 1$ and you receive \$10,000.

Plausibly, for any such one-off choice, you are required to press the button. After all, the difference in pain is tiny and you get \$10,000! But now consider:

Torture

You must choose between receiving no shocks and no money or receiving shocks at level 1,000 and \$10,000,000.

Plausibly, if the pain is bad enough, you are required to choose the painless option.

If one accepts both judgments, then DP 3 leads to contradiction. Imagine facing a sequence of 1,000 choices as described in Extra Turn. By ME 3 and Extra Turn, you are required to press each time. Then, by ICS 3, you are permitted to choose the sequence

¹¹ This case comes from Quinn 1990, but some doubt the permissibility of cyclical preferences even in such cases. See Arntzenius 1997 for an influential critique and Tenenbaum 2012 for a reply.

(press the button + ... + press the button). Then, by SO 3, you are permitted to choose the outcome of (press the button + ... + press the button) over the outcome of (don't press + ... + don't press). So you are permitted to choose shocks at level 1,000 and \$10,000,000 over no shocks and no money. This contradicts Torture.

Section 3.4: A Common Assumption

The DPs conflict with many intuitive judgments. Opponents of PA might, in response, endorse only some of the DPs and reject the others. For example, they might endorse only DP 3 and thus accept parity, supererogation and prerogatives, while rejecting PA and cyclical preferences. Since cyclical preferences are far more controversial than PA, the opponents of PA could happily reject both.

Such views face a problem: why accept one of the DPs while rejecting the others? Without a good explanation, the view looks ad hoc. But it seems hard to imagine either a principled divide among the DPs, or reasons to accept some of the DPs that do not generalize.

Furthermore, there is an attractive diagnosis of what goes wrong with the DPs that applies to all three principles: they rely on a common but false assumption about diachronic choice.

To see this, start by escaping the weeds. In general, there are two kinds of theories about how agents should choose over time. One kind requires agents to focus just on the choice presently in front of them, asking them to evaluate how the effects of that choice alone compare (including effects on what actions may be available later). Call this the Immediate Perspective. Another requires agents to consider all the outcomes that could

arise from any available sequence of choices and to pick a sequence that results in one of the best outcomes. Call this the Planning Perspective.

It may help to consider an example. Each afternoon, I can either work on my papers or get coffee with a friend. The Immediate Perspective tells me to, each afternoon, compare my reasons for working that afternoon (e.g. improvements in the prose) with my reasons for getting coffee with a friend that afternoon (e.g. delightful conversation). These reasons might also include (i) effects which take place only after a long time, and (ii) effects on which options I may have later. So, for example, if I chose to work on my friend's birthday, she may eventually come to resent this and henceforth refuse to get coffee with me. The Immediate Perspective can say that, on my friend's birthday, both the future resentment and reduction of my future options are reasons to choose coffee. Meanwhile, the Planning Perspective tells me to consider all the possible outcomes available to me (e.g. the outcome of work on day 1 + coffee on day 2 + ... + work on day 83 + coffee on day 84) and to choose a sequence that produces one of the best outcomes.

The conjuncts of DP, namely Moral Equivalence, Sequences to Outcomes, and Individual Choice to Sequences, naturally align with one or another perspective.

Moral Equivalence states, roughly, that if two sets of actions, across different decision problems, offer the same effects (including effects on what actions will be available in the future), then one action is permitted/required iff its counterpart is permitted/required. This thought makes sense from the Immediate Perspective, where one assesses each action on its own merits.

Sequences to Outcomes states, roughly, that choosing one sequence of actions over another is just like choosing the former outcome over the latter. That thought makes sense from the Planning Perspective, where one cares about the whole sequence of actions.

By contrast, Individual Choice to Sequences, fits with both perspectives. It is a kind of consistency condition, capturing the thought that morality speaks with one voice. Morality would not be so silly, for example, as to require you to perform each X_i , and yet fail to require the whole sequence ($X_1 + \dots + X_n$). So if each individual action was required, then the sequence must be required (as ICS 2 asserts).

Now, for any particular decision problem, there are three possibilities:

Option 1: The Immediate Perspective alone governs how one should choose.

Option 2: The Planning Perspective alone governs how one should choose.

Option 3: Both the Immediate Perspective and the Planning Perspective govern how one should choose.

Among these, only Option 3 justifies imposing the DPs. So the DPs share a common assumption: that Option 3 applies to the case at hand.

Seeing this common commitment gives reason to doubt the DPs since, prima facie, Option 3 is rather strange. The Immediate Perspective and the Planning Perspective are very different from one another, and as earlier cases show, combining these perspectives with plausible assumptions can yield contradictions. These perspectives are at odds with one another. So why should anyone expect that, in any particular case, they both govern?

Perhaps it is simply obvious that, in ordinary cases, Option 3 applies. After all, the Immediate Perspective is an extremely natural way to think about ordinary choices, and so

is the Planning Perspective. The best explanation is that both perspectives govern ordinary choices. So typically, Option 3 is true; the above cases are just aberrations.

But there is another explanation of these intuitions: in ordinary (i.e. finite, transitive) settings, the prescriptions of the Immediate Perspective and the Planning Perspective coincide.¹² For suppose, to modify the preceding example, I prefer the outcome of (work + coffee) to (coffee + work) to (work + work) to (coffee + coffee). I care most about balancing work with pleasure, and I prefer, *ceteris paribus*, to get my work done first. Then, the Planning Perspective tells me to choose the sequence (work + coffee). And the Immediate Perspective tells me to first choose work (since I prefer to get my work done first), and then to choose the coffee (for work-life balance). The two perspectives coincide in their prescriptions.

This clears the path for an error theory. In some cases, the Immediate Perspective alone governs, and in others, the Planning Perspective alone governs. Which is true in any particular case may depend on substantive considerations relevant to the case at hand.

Perhaps, for example, Islands is governed only by the Immediate Perspective, while an earlier example of parity, Case 7, is governed only by the Planning Perspective. After all, in Islands, your actions affect each person once, whereas in Case 7, your actions affect the same person multiple times. The separateness of persons may thus prohibit agents from mixing reasons across choices in Islands, while allowing agents to mix reasons across choices

¹² I have not characterized either perspective with enough precision to prove this. But if I formalized the Immediate Perspective as naive or sophisticated decision theory, and the Planning Perspective as resolute decision theory, then I could prove: If the pairwise choice relation is transitive and there are finitely many nodes, then both perspectives permit the same paths through any decision tree. (For the relevant formal definitions, see McClennen 1990 and Cubitt 1996).

in Case 7. So the Immediate Perspective applies to the former, while the Planning Perspective applies to the latter.

In ordinary (i.e. transitive) cases, such complications do not matter in determining what one ought to do. As a result, one's intuitions are not attentive to which perspective governs. Since both perspectives are plausible, it then seems like both perspectives govern in ordinary cases. But in fact, ordinary intuitions alone do not support this conclusion; they do not support Option 3 over (Option 1 v Option 2).

Section 4: Explaining PA's Violations of the DPs

I have given some general reasons to doubt the DPs. But proponents of PA can also provide explanations for why the DPs should fail in cases of PA. This section provides one account that combines PA with the Immediate Perspective and sketches another that combines PA with the Planning Perspective.

Section 4.1: The Immediate Perspective

The Immediate Perspective, when applied to Islands, requires the agent to choose the sequence of (save D + save E + save F), even though, in a one-off choice, one would be required to save the denizens of A+B+C over the denizens of D+E+F.¹³ This combination of verdicts may appear bizarre, but it can be explained as follows.

Return to Monday on Islands. You see island A to your left and island D to your right.

You can only save one. What reasons bear on which island you ought to save?

¹³ I address the case of Islands, rather than Horton's original cases, because cases of risk include additional complications that are not needed for Horton's diachronic challenge since the DPs suffice to rule out PA.

Plausibly, your only reasons for action are those provided by the potential suffering of the islanders on A and D. For clearly their potential suffering does provide you with reasons to act. And plausibly, the islanders on, say, F do not provide you with a reason to act on Monday, when forced to choose between A and D. After all, nothing you do on Monday will affect F, nor will it affect whether you can affect the islanders on F. So how could the F's provide you with a reason to act on Monday? Finally, if your only reasons for acting on Monday are given by A and D, then D's reasons carry the day. For in a one-off choice between saving A and D, you must save D.

To assess these claims, it may help to consider another case:

Attenborough 1

No matter what I do, David Attenborough will have a wonderful Sunday lunch.

Intuitively, David Attenborough's fate does not provide me with any reasons to act. Of course, this case's structure differs from Islands. So consider:

Attenborough 2

Suppose that I have a strange power: next Sunday, if I snap, then David Attenborough's lunch will become far tastier. Still, today is Monday, and nothing I do in the interim will either affect David Attenborough or whether I continue to wield this awesome power.

Intuitively, David Attenborough's fate does not provide me with any reasons to act between now and Sunday.

If you agree with these verdicts, then you should make the corresponding judgment about Islands: on Monday, only A and D's fates provide reasons for action. Thus, on Monday, you should save D. Similar remarks apply to the other days.

By contrast, in a one-off choice between A+B+C and D+E+F, everybody is affected and hence, everybody's fate yields a reason for action. And it is plausible that the reasons for saving A+B+C outweigh those for saving D+E+F. So, one must save A+B+C.

This pattern of verdicts, far from being bizarre or ad hoc, falls out of PA in combination with two natural thoughts about how reasons operate: (i) what you ought to do at some choice in a sequence depends only on what reasons you have then, and (ii) if your action will neither affect X, nor affect how you can affect X, then ceteris paribus, X's fate does not yield a reason for action.

Section 4.2: The Planning Perspective

The Planning Perspective, when applied to Islands, requires the agent to choose the sequence of (save A + save B + save C), even though, in a one-off choice, one would be required to save D over A, to save E over B, and to save F over C. Again, these verdicts can be motivated by plausible general principles about reasons.

Start with the following thought: sometimes, X's fate can give you reasons to act even though your actions neither affect X nor affect how you can affect X later. Consider:

Attenborough 3

Just like Attenborough 2, except now your friend, who hates nature documentaries, asks you to promise (on Monday) not to snap your fingers next Sunday.

Your promise will neither affect Attenborough, nor affect whether you can affect him on Sunday (your awesome powers remain regardless). But intuitively, Attenborough's fate does give you some reason not to promise. One diagnosis of this intuition would be:

Reasons for Reasons

If, by ϕ -ing, you will add to (detract from) the reasons you later have for ψ -ing, and X's fate gives you reason to ψ , then X's fate also gives you reason to (not) ϕ .

Perhaps some principle like this is true. And perhaps, in Islands, whether one saves A or D on Monday will affect the reasons one has, on subsequent days, for saving one or another group. Then, perhaps one has additional reasons to save A rather than D in Islands which are not present in an otherwise similar one-off case.

Such thoughts point to avenues for developing a principled explanation for why PA should reject the DPs and endorse the Planning Perspective alone.

Section 4.3: Horton's Objections

Horton considers and objects to both views. But his objection to the Immediate Perspective appears to support the Planning Perspective, and vice-versa. So it helps to look at these objections together.

Horton's objection to the Immediate Perspective goes as follows.¹⁴ Consider:

Islands*

Just like Islands except, on Monday, you have two additional options: save A+B+C all at once (condemning the rest), or save D+E+F all at once (condemning the rest).

In Islands, the Immediate Perspective requires the agent to first save D, then E, and then F. But plausibly, in Islands*, the Immediate Perspective requires the agent to simply save A+B+C on Monday. It is absurd, though, to think that one can be required to make a certain series of choices while being prohibited from simply skipping ahead to the final outcome. Could it really be so morally significant that someone has to choose on three separate days?

¹⁴ Horton 2020, pg. 520. I have modified both objections to fit Islands.

This objection has great intuitive force. But it relies on the thought that, when asked to make a series of choices over time, the “internal structure” of the choice problem cannot matter. When exactly one must make certain choices, and between which options, is simply irrelevant (except, perhaps, as it may indirectly bear on the intentions or motives of the agent). What really matters is which final outcome results. Such thoughts, though, are exactly what motivate the Planning Perspective.

Now consider Horton’s objection to the Planning Perspective:¹⁵

Islands**

You know that you *may* face a series of choices like Islands. But your heart is weak and you are unsure whether you will face all three choices, or die after making the first choice. Today is Monday.

If you face one choice, then the Planning Perspective, like all PA views, says to save D. If you face all three choices, then the Planning Perspective says to save A, save B, and then save C. But intuitively, it is absurd to think that whether your heart will give out could alter what one ought to do.

Such intuitions are forceful. But they rely on the thought that your reasons for acting on Monday, even if you must make later choices, are given just by the participants affected by your Monday actions. That is why it cannot matter whether one lives long enough to make choices on Tuesday and Wednesday. Such thoughts motivate the Immediate Perspective.

So it is true that Horton offers forceful objections to both perspectives. But each objection motivates the other perspective. So one should respond, not by rejecting PA, but

¹⁵ Horton 2020, pg. 519.

by making up one's mind about whether the Immediate Perspective or the Planning Perspective governs in these cases. After doing so, Horton's objections may lose their force.

Conclusion

I have shown how Horton's arguments can be strengthened, so as to avoid reliance on risky cases. Those arguments force PA to choose between two ways of looking at choices over time: the Immediate Perspective and the Planning Perspective. PA cannot accept both perspectives, but this fact alone does not rule out PA since many other normative phenomena display the same pattern.

Furthermore, every normative theory must take some stand on the choice between the Immediate and the Planning Perspective. For every normative theory has two components: (i) a theory of weights, stating how reasons interact with one another, and (ii) a theory of reasons, stating which reasons (and in virtue of which facts) apply to any particular choice. So any normative theory must answer whether, on Monday in Islands, B/C/E/F yield reasons for action, or not.

Other (i.e. transitive) normative theories are lucky in that both perspectives will yield the same verdicts. If we only ask about what agents ought to do, such views will overlook this issue. But strictly speaking, they must also take a stand. And if we can agree on a general theory of reasons, and thus choose between the Immediate and Planning Perspective, PA does not appear to face a further problem.

Bibliography

- Arntzenius, Frank and David McCarthy (1997). Self-Torture and Group Beneficence. *Erkenntnis* 47 (1): 129-144.
- Bader, Ralf (2019). Agent-Relative Prerogatives and Sub-Optimal Beneficence. *Oxford Studies in Normative Ethics* 9: 223-250.
- Chang, Ruth (1997). Introduction. *Incommensurability, Incomparability, and Practical Reason*. Harvard University Press: 1-34.
- Chang, Ruth (2002). The Possibility of Parity. *Ethics* 112 (4):659-688.
- Chang, Ruth (2005). Parity, Interval Value, and Choice. *Ethics* 115 (2):331-350.
- Cubitt, Robin (1996). Rational Dynamic Choice and Expected Utility Theory. *Oxford Economic Papers* 48 (1): 1-19.
- Dougherty, Tom (2014). A Deluxe Money Pump. *Thought: A Journal of Philosophy* 3 (1):21-29.
- Gustafsson, Johan & Wlodek Rabinowicz (2020). A Simpler, More Compelling Money Pump with Foresight. *Journal of Philosophy* 117 (10):578-589.
- Hare, Caspar (2010). Take the Sugar. *Analysis* 70 (2):237-247.
- Horton, Joe (2018). Always Aggregate. *Philosophy and Public Affairs* 46 (2):160-174.
- Horton, Joe (2020). Aggregation, Risk, and Reductio. *Ethics* 130 (4):514-529.
- McClennen, Edward. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press.
- Muñoz, Daniel (2021). Three Paradoxes of Supererogation. *Noûs* 55 (3):699-716.
- Parfit, Derek (2011). *On What Matters: Volume One*. Oxford University Press.
- Quinn, Warren S. (1990). The Puzzle of the Self-Torturer. *Philosophical Studies* 59 (1):79-90.
- Tenenbaum, Sergio & Diana Raffman (2012). Vague Projects and the Puzzle of the Self-

Torturer. *Ethics* 123 (1):86-112.