

Agent-relative prerogatives and suboptimal beneficence

Ralf M. Bader
Merton College, University of Oxford

I Suboptimal beneficence

Agent x has two options:

- A = do nothing
- B = save y's left arm at some cost to himself

Both A and B are permissible. Given that the cost $-\alpha$ that x has to incur to save y's arm is sufficiently large relative to the benefit $+\beta$ conferred on y, x is permitted but not required to benefit y. Helping y is good but not required and hence classifies as supererogatory.

- A = do nothing
- C = save both of y's arms at the same cost to himself

Although the benefit is greater than in the previous case, helping y is not required if the cost to x is sufficiently large. Both A and C are then permissible.

- A = do nothing
- B = save y's left arm at some cost to himself
- C = save both of y's arms at the same cost to himself

When all three options are available, however, only A and C are permissible. B is impermissible.¹ As Parfit notes, saving only one arm when one can save both at the same cost is 'grossly perverse' (cf. Parfit: 1982, p. 131 and Kagan: 1989, pp. 15-17 & 240-241). Suboptimal beneficence of this kind is impermissible. Consequently, neither supererogation nor agent-relative prerogatives can be understood

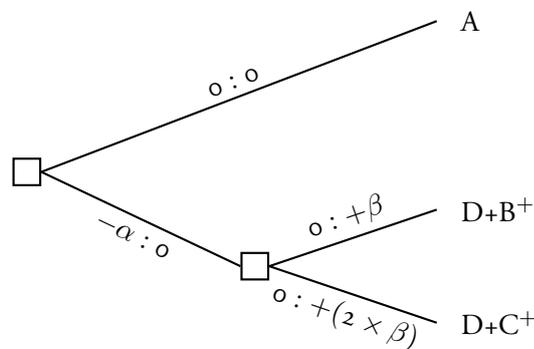
¹Taking the default to be A, we can represent the choice situation as follows (n.b. nothing hangs on construing A to be the default – this is simply done for illustrative purposes).

	x	y
A	0	0
B	$-\alpha$	$+\beta$
C	$-\alpha$	$+(2 \times \beta)$

in terms of threshold models.² Such models include options that are above the threshold, but are nevertheless suboptimal in this objectionable way and hence impermissible. In addition, the impermissibility of suboptimal beneficence is of significant practical relevance.³

2 Sequential choice

The intuitive verdict that only A and C are permissible when choosing from the triple {A, B, C} receives further support from sequential cases. An agent first decides whether to do nothing = A or to incur the cost = D. If the cost is incurred, he faces a subsequent choice between two (by then costless) beneficent actions: saving only y's left arm = B⁺ or saving both arms = C⁺.



The following three diachronic choice consistency conditions, namely (i) separability from choice history, (ii) dynamic consistency, and (iii) reduction of sequential choice,⁴ together with the substantive claim that (iv) in a choice between costless versions of B and C, namely B⁺ and C⁺, where doing nothing is not an option, only C⁺ is permissible, imply that in the nonsequential case only A and C are permissible whereas B is impermissible.⁵

First, separability implies that the choice between B⁺ and C⁺ as part of the larger decision tree has to be the same as that between these options in an isolated tree. The agent has to choose at node n_1 in the subtree in the same way as at node n_0 in a corresponding isolated tree.⁶

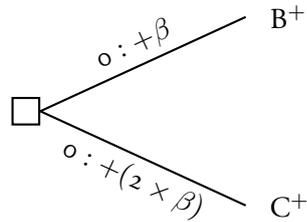
²Nor in terms of moving bands (cf. Hurka and Shubert: 2012).

³Cf. Pummer: 2016 who has argued that these cases support effective altruism.

⁴Cf. Cubitt: 1996 for formal statements of these conditions.

⁵The same reasoning applies when the agent first decides whether to incur the cost in order to save one arm and, if the cost is incurred, then faces a subsequent decision whether to save the second arm at no additional cost

⁶This argument only requires separability from choice history, not separability from chance history. The latter applies when deciding under uncertainty and is rather contentious, given its connection to the independence axiom of expected utility theory. Cf. Cubitt: 1996 for a precise characterisation of these separability principles.

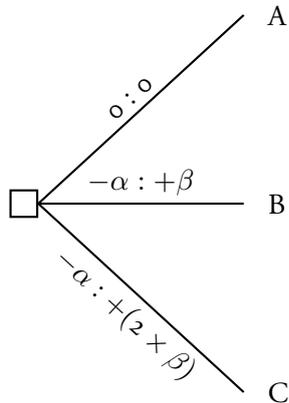


This choice is between two costless acts of beneficence. (Given a counterfactual baseline for evaluating benefits, only C^+ will classify as an act of beneficence.) There is no reason (and no excuse) for choosing the suboptimal one. B^+ is impermissible. Separability then implies that B^+ is likewise impermissible in the subtree at node n_1 . First incurring the cost and then performing the suboptimal beneficent action is as problematic as choosing B^+ when deciding between two costless acts of beneficence. The fact that a cost was previously incurred that could have been avoided by doing nothing does not affect the permissibility of this choice.⁷

Second, dynamic consistency requires that choices at subsequent nodes are consistent with the strategies adopted at prior nodes. Plans adopted at later nodes have to be continuations of the original plan. Subsequent evaluations have to agree with the initial evaluation. By requiring the choice at n_1 to be a continuation of a permissible strategy at n_0 , dynamic consistency establishes that the only strategy that goes via n_1 that is permissible from the outset is the one that results in C^+ , given that one has to choose C^+ rather than B^+ at the subsequent choice node.

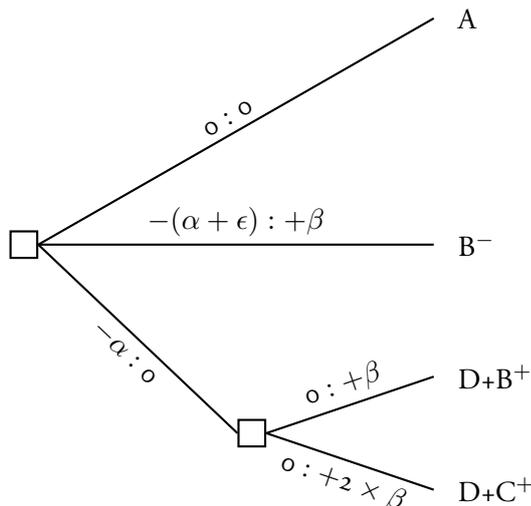
Third, reduction of sequential choice requires the sequential case to be treated analogously to the nonsequential case. Sequential choice has to agree with nonsequential choice. The permissible strategies in the complex tree have to correspond to those that are permissible in the nonsequential tree that has the same normal form.

⁷Some complications arise. We will see in section 7 that a commitment to agent-relative prerogatives is only compatible with dynamic consistency when using global choice principles. Global choice, however, is incompatible with separability, which seems to undermine the first step of the argument. Yet, we can specify a restricted separability principle that is compatible with global choice and that allows us to establish the impermissibility of suboptimal beneficence.

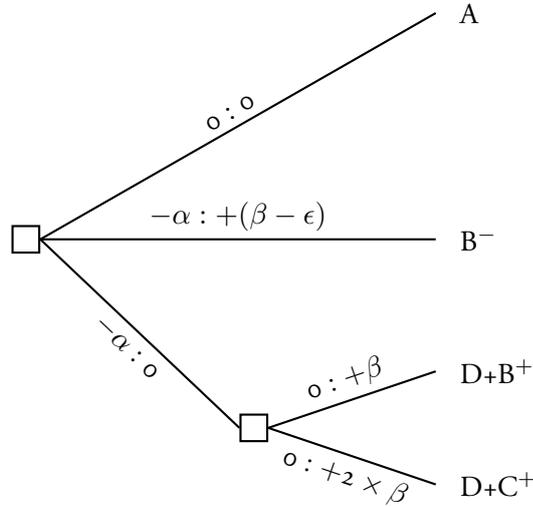


Since $\{A, B, C\}$ and $\{A, D+B^+, D+C^+\}$ have the same payoffs, reduction implies that B is permissible iff $D+B^+$ is permissible. Since B^+ is impermissible (given separability), the strategy $D+B^+$ is impermissible at n_0 (given dynamic consistency), which implies that B is impermissible in the nonsequential case (given reduction). Suboptimal beneficence is, accordingly, impermissible, even when doing nothing is permissible, both in the sequential and the nonsequential case.

Those, such as McMahan: 2017, who consider all three options to be permissible when choosing nonsequentially from $\{A, B, C\}$, yet consider the equivalent of B to be impermissible when choosing sequentially from $\{A, D+B^+, D+C^+\}$, have to reject dynamic consistency or reduction of sequential choice. It will then be permissible in certain cases to ‘pay’ (either by incurring a greater cost, or by foregoing to confer a greater benefit) in order to achieve an outcome that is available but that otherwise could not permissibly be brought about. The agent can permissibly reach the outcome of B^- but not the superior outcome of $D+B^+$.



Here one is allowed to bring about benefit β but only at a larger cost to the agent, namely $-(\alpha + \epsilon)$, but not at the smaller cost of only $-\alpha$.



Here one is allowed to incur cost $-\alpha$ in order to confer benefit $+(\beta - \epsilon)$ but not allowed to incur this cost in order to confer a slightly greater benefit $+\beta$.

3 Agent-relative prerogatives

The prohibition on suboptimal beneficence follows from an agent-relative theory that understands permissible actions in terms of a dominance principle defined over both the agent-relative and the agent-neutral ordering. This theory incorporates agent-relative prerogatives, allowing agents to act on the basis of the agent-relative evaluation that is centred on the agent. The agent can permissibly privilege (i.e. be partial towards) himself as well as those who are close to the agent and is not required to be impartial and do what is agent-neutrally best.

An agent-relative axiological theory consists of an agent-relative and an agent-neutral ordering. Alternatives can be evaluated in two different ways. The evaluation is either indexed on the perspective of the agent, in which case it is sensitive to identity facts and thereby allows for partiality. Or it abstracts from the identity of the agent, making it an impartial evaluation that is permutation-invariant.

These two evaluations are connected. The agent-relative evaluation results when applying a modification function to the agent-neutral evaluation.

$$V_{\text{relative}}(y) = k \times V_{\text{neutral}}(y), \text{ where } k = f(d(x, y))$$

The modifier k is a function f of distance d in moral space between the agent x on whom the evaluation is centred and the person y whose well-being is evaluated, where distance itself is a function of morally significant relationships between the persons. The strength of the multiplier is inversely related to distance. The greater the distance, the smaller the multiplier. The closer someone is to the agent, the more that person counts from the agent's perspective. For a complete stranger

$k = 1$, i.e. $f(d(x, \text{stranger})) = 1$, which means that the agent-relative reason is equally strong as the agent-neutral reason, whereas $\max k = f(d(x, x))$, i.e. no one is more closely related to the agent than the agent himself.⁸

We have two standpoints from which things can be evaluated, each giving rise to a complete ordering that weighs up all the costs and benefits: the agent-relative and the agent-neutral standpoint. Although the two evaluations are connected via the modification function k , the resulting orderings cannot be weighed up against each other (cf. Bader: 2016, section 4.1 and Bader: 2015).⁹ This is not merely an epistemic concern about our ability to precisely weigh up the different values. Nor is it an issue of incommensurability, given that the modification function ensures that the two evaluations are commensurable. Instead, the issue is that the values are internal to their respective standpoints and that there is no further standpoint that subsumes them both. There is no further perspective that allows us to combine these evaluations by determining their relative weights. All we have are these two standpoints, neither of which is privileged over the other.¹⁰

The top-elements of the two orderings are endorsed by their respective standpoints. Since neither standpoints is subordinated, in particular since neither standpoint is privileged over the other, there is nothing that can counteract their respective endorsements. The agent can hence permissibly perform those actions that are undominated on at least one ordering. He can either do what is agent-relatively best or what is agent-neutrally best. The following dominance principles defined over both orderings yield the set of permissible actions.¹¹

WEAK PURE DOMINANCE

$$C(X) = \{x \in X : \neg \exists y \neg \exists z \in X (y >_{\text{relative}} x \wedge z >_{\text{neutral}} x) \wedge \neg \exists w \in X [(w \geq_{\text{relative}} x \wedge w >_{\text{neutral}} x) \vee (w >_{\text{relative}} x \wedge w \geq_{\text{neutral}} x)]\}$$

STRICT PURE DOMINANCE

⁸For an account of modification functions, cf. Bader: 2016. For a detailed treatment of this theory of agent-relativity cf. “Personal, general, and impersonal good” Bader: manuscript.

⁹There would be no room for prerogatives if there were to be a correct way of weighing them up. Cf. Hurka and Shubert: 2012, pp. 3-5 who argue that permissions cannot be generated when operating with values or reasons and hence advocate positing underivative permissions.

¹⁰Given that there is not one way to weigh them up, one might consider all possible ways of weighing them up and taking an action to be permissible if it is permissible according to some way of weighing them up. This amounts to taking the union of the optimising choice sets of all convex combinations of the two evaluations, i.e. $V_{\text{combined}} = \lambda V_{\text{relative}} + (1 - \lambda) V_{\text{neutral}}$. The difference between λ being in the closed or open unit interval corresponds to the difference between strict and weak dominance principles: strict if $\lambda \in [0, 1]$ and weak if $\lambda \in (0, 1)$. (The two asymmetric principles correspond to $\lambda \in (0, 1]$ and $\lambda \in [0, 1)$.) This way of weighing up the two evaluations, however, corresponds neither to the pure nor the mixed dominance principles. It is more permissive than the pure principle and less permissive than the mixed principle. Moreover, it does not satisfy dynamic consistency.

¹¹These principles are somewhat similar to the dual-ranking theories developed by Sider: 1993 and Portmore: 2008, which are based on the self v. other distinction rather than the agent-relative v. agent-neutral distinction. (Thanks to Doug Portmore for pointing out these articles to me.)

$$C(X) = \{x \in X : \neg \exists y \neg \exists z \in X (y >_{\text{relative}} x \wedge z >_{\text{neutral}} x)\}$$

The strict dominance principle rules out those alternatives that are not top-elements on either ordering. The weak principle also rules out alternatives that are tied for top-element on one ordering with something that dominates them on the other ordering. The weak principle treats the agent-relative and the agent-neutral orderings symmetrically. There are also two asymmetric principles that require only one of the disjuncts featuring in the weak dominance principle: either $\neg \exists w \in X (w \geq_{\text{relative}} x \wedge w >_{\text{neutral}} x)$ or $\neg \exists w \in X (w >_{\text{relative}} x \wedge w \geq_{\text{neutral}} x)$. In order to argue in favour of the weak principle (over the strict principle as well as over the two asymmetric principles), two possibilities need to be ruled out.

First, is one allowed to perform an action that is agent-relatively worse yet tied in terms of agent-neutral value? Here the question is whether one can set aside agent-relative reasons and take up an agent-neutral perspective, even when this does not make things impartially better. For instance, is it permitted to help a stranger rather than a friend when the benefit is the same, so that these options are equally good from the agent-neutral perspective, yet the latter is preferable from the agent-relative perspective? By citing the importance of friendship one can justify to the stranger that one is not helping him but instead one's friend. Yet it would seem that one cannot justify to one's friend why one is not helping him. Helping the stranger could be justified if doing so would bring about more agent-neutral good. However, when the alternative is not impartially better, then this justification is not available.

Second, is one allowed to perform an action that is suboptimal agent-neutrally but tied in terms of agent-relative value? Performing such an action looks like a gratuitous failure to promote the impartial good. The alternative effectively amounts to a case of 'costless' beneficence, insofar as there is no net cost in terms of agent-relative good (though it may involve costs to the agent), and hence would seem to be required. Given that the two actions are equally good agent-relatively, one cannot cite any agent-relative considerations to justify not performing that action which is agent-neutrally optimal.

These considerations suggest that we should adopt the weak rather than strict dominance principle. The arguments of this paper, however, go through on both versions. In each case the dominance principle can account for the permissibility pattern with which we started.

When choosing between A and B, both are permissible. The agent-relative ordering is: $A > B$. The agent-neutral ordering is: $B > A$. This ensures that neither option is dominated, so that both options end up in the choice set. The same reasoning explains why both A and C are permissible if one is to choose between them. When choosing between all three alternatives, however, the agent-relative ordering is: $A > C > B$, whereas the agent-neutral ordering is: $C > B > A$. The dominance principles allow the agent to either pick option A (which is agent-relatively best) or option C (which is agent-neutrally best). Options, such as B,

that are suboptimal on both orderings, by contrast, are impermissible. The agent-relative prerogative justifies x in choosing A rather than C, but does not justify choosing B rather than C, thereby ensuring that A and C end up in the choice set whereas B is excluded.

4 Mixed dominance

The reasoning that renders B impermissible when choosing from $\{A, B, C\}$ also applies when the two different ways of helping do not involve the same cost, but when the agent needs to incur a slightly larger cost to confer the greater benefit.

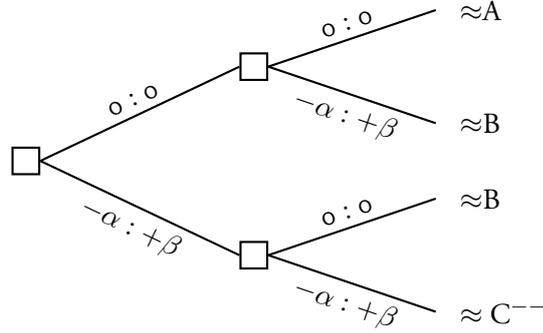
	x	y
A	o	o
B	$-\alpha$	$+\beta$
C^-	$-(\alpha + \epsilon)$	$+(2 \times \beta)$

A and C^- are the only permissible options as long as ϵ is sufficiently small. B does become permissible when the additional cost is sufficiently large. For instance, B is no longer impermissible if the additional cost is as large as the original cost, i.e. when C^{--} involves a cost of $-(2 \times \alpha)$ for x and a benefit of $+(2 \times \beta)$ for y .

	x	y
A	o	o
B	$-\alpha$	$+\beta$
C^{--}	$-(2 \times \alpha)$	$+(2 \times \beta)$

The pure dominance principles, however, exclude B from the choice set since B is suboptimal on both orderings. Yet, intuitively, B should be permissible when C^{--} is the alternative. This becomes clear when thinking of an analogous sequential choice situation, whereby x is faced on two separate occasions with the option of doing nothing or conferring benefit β at cost α .

	x	y
\approx A	o, o	o, o
\approx B	$o, -\alpha$	$o, +\beta$
\approx B	$-\alpha, o$	$+\beta, o$
\approx C^{--}	$-\alpha, -\alpha$	$+\beta, +\beta$



Intuitively, both alternatives are permissible at each node. That x has (not) helped at the first choice node does not imply that he has to (not) help again at the second node. Although we will reject separability in section 7, the reasons for doing so do not apply to this case, given that the initial choice makes no difference to the rest of the decision tree. The sub-tree that one faces at the second node is the same independently of whether one has helped at the initial node. The choice at n_0 has no effect on what options are available at n_1 : in each case one can either do nothing or confer benefit β at cost α . Moreover, if separability were to fail in this case, it would support rather than rule out mixed strategies, since it is understandable that an agent does not help again if he has already helped in the past. It is, accordingly, permissible for the agent to help only once. Since helping once is the analogue of choosing B from $\{A, B, C^{--}\}$, reduction of sequential choice implies that B is likewise permissible.

A more nuanced approach is thus required. The pure dominance principles allow one to either adopt the relative or the neutral perspective. The agent can then perform actions that are permissible according to at least one point of view. We can weaken these principles and allow the agent to perform any action that is not dominated by some alternative on both orderings.¹²

WEAK MIXED DOMINANCE

$$C(X) = \{x \in X : \neg \exists y \in X [(y \geq_{\text{relative}} x \wedge y >_{\text{neutral}} x) \vee (y >_{\text{relative}} x \wedge y \geq_{\text{neutral}} x)]\}$$

STRICT MIXED DOMINANCE

$$C(X) = \{x \in X : \neg \exists y \in X (y >_{\text{relative}} x \wedge y >_{\text{neutral}} x)\}$$

If some alternative dominates an action from both the agent-relative and the agent-neutral perspective, then that action is impermissible. By contrast, if both orderings agree that an action is suboptimal but they disagree on the reasons for which it is suboptimal, then it is not impermissible. To be impermissible, the

¹²The principles are weaker in that they are more permissive. It is easier for things not to be dominated since the condition being specified by the dominance principle is stricter such that less things can satisfy this condition.

orderings not only have to agree that something is dominated, i.e. that it is not a top element, but also have to classify it as being suboptimal for the same reason, i.e. they have to agree on which alternative dominates it. Unless they agree, the reason for which the one perspective rules out a given action will not be admissible from the standpoint of the other, and vice versa. There will then be no agreed reason for ruling it out.

To be responsive to the fact that an action is suboptimal is to not perform this action but instead one that dominates it. If there is no agreed reason, the agent cannot be responsive to both perspectives. There will be an action that dominates it on the agent-relative ordering and an action that dominates it on the agent-neutral ordering, but no action that dominates it on both orderings. The agent can then either be responsive to the agent-relative or the agent-neutral perspective, but not to both. Even though the other point of view agrees that the action is to be ruled out, the agent cannot be responsive to both of the reasons for which it is ruled out. Instead, the agent will end up siding with one perspective and will do something that the opposing point of view deems to be worse than the action that both points of view consider to be suboptimal. It is for this reason that one should operate with the mixed dominance principles and consider actions to be permissible, even if they are suboptimal on both orderings, as long as there is no agreed reason for considering them to be suboptimal.

These mixed dominance principles ensure that, whilst B is impermissible when A and C^- are the other alternatives, B is permissible in the case of C^{--} . As long as the additional cost is sufficiently small, C^- is above B in both rankings. The cost has to be so small that the agent would have to incur $-\epsilon$ to confer benefit $+\beta$. Helping is not optional but required when the value of the additional benefit is greater than the additional cost, even from the agent-relative perspective, i.e. $k \times \epsilon \leq \beta$. (We need a strict rather than weak inequality if we are using the weak rather than strict mixed dominance principle.) However, if the additional cost is sufficiently large, i.e. if $k \times \epsilon > \beta$, then it is optional and not required to incur this cost to confer benefit β .¹³ The mixed principle then renders B permissible. Whilst B is ranked below C^- on both orderings (relative: $A > C^- > B$ and neutral: $C^- > B > A$), B is ranked above C^{--} on the agent-relative ordering (relative: $A > B > C^{--}$ and neutral: $C^{--} > B > A$). There is hence no option that dominates B on both orderings. Whilst being dominated on both orderings, B is dominated by one option, namely C^{--} , on the agent-neutral ordering, yet by another option, namely A, on the agent-relative ordering.

The mixed dominance principles which require the different orderings to agree allow us to construct a single combined evaluation, namely an intersection quasi-ordering. Since it is an incomplete ordering, we have to apply a maximis-

¹³The modification function, in this way, determines precisely how much greater to x the additional cost has to be in order for B to be permissible.

ing rather than optimising function to determine the choice set.¹⁴ The difference between the weak and the strict mixed dominance principles over the two orderings concerns the construction of this quasi-ordering. The difference is whether $x \gg y$ only requires x to be strictly better than y with respect to one of $>_{\text{relative}}$ and $>_{\text{neutral}}$ whilst being weakly better with respect to the other, or whether x needs to be strictly better on both orderings.

WEAK INTERSECTION

$$x \gg y =_{\text{df}} x \geq_{\text{relative}} y \wedge x \geq_{\text{neutral}} y \wedge \neg(x \equiv y)$$

where

$$x \equiv y =_{\text{df}} x =_{\text{relative}} y \wedge x =_{\text{neutral}} y$$

STRICT INTERSECTION

$$x \gg y =_{\text{df}} x >_{\text{relative}} y \wedge x >_{\text{neutral}} y$$

The choice set then contains all alternatives that are not strictly dominated with respect to the intersection quasi-ordering.

$$C(X) = \{x \in X : \neg \exists y \in X (y \gg x)\}$$

5 Prerogatives and beneficence

This account explains not only why it is impermissible to help suboptimally, yet permissible to not help. It also explains why it is impermissible to help suboptimally, even though it is permissible to not do what is impartially best but instead help those who are close to the agent. A moderate agent-relative theory that uses a modifier k that is a continuous function of distance in moral space gives rise to an agent-relative prerogative that not only allows the agent to privilege his own interests but also the interests of those that are close to him.¹⁵ Accordingly, one can help a relative rather than a stranger, even if the latter would benefit more. Yet one is not permitted to engage in suboptimal beneficence either with respect to strangers or with respect to one's relatives.

	x	y	r
A*	o	o	$+\gamma$
B*	o	o	$+(2 \times \gamma)$
C*	o	$+\beta$	o
D*	o	$+(2 \times \beta)$	o

¹⁴An optimising choice function, which selects options that are at least as good as all the alternatives, presupposes completeness. A maximising choice function, which selects undominated options, is compatible with incompleteness (cf. Sen: 1997).

¹⁵The modification function k does not only apply to the agent but to everyone.

Since the modifier varies with distance in moral space, greater weight is assigned to the well-being of those who are closer to the agent than those who are more distant. If the relative r is sufficiently close to the agent x , then the modifier $k = f(d(x, r))$ can ensure that, even though $(2 \times \gamma) < \beta$, it is nevertheless the case that $k \times \gamma > 2 \times \beta$ (where y is a stranger so that the multiplier = 1). The agent is then permitted to take up either the agent-neutral perspective and perform D^* or the agent-relative perspective and perform B^* . A^* and C^* , however, are impermissible. They are suboptimal on both orderings. Agent-relative as well as agent-neutral suboptimal beneficence are ruled out. We thus have a unified account that encompasses both partiality to oneself and to others.¹⁶

The prohibition of this kind of suboptimal beneficence, by contrast, cannot be explained in terms of conditional obligations. Whilst one might be able to explain the impermissibility of B when choosing from $\{A, B, C\}$ by arguing that, even though one is not required to help, one is obligated to help in the optimal way on condition that one does help, this approach cannot account for the prohibition of suboptimal beneficence (i.e. the impermissibility of A^* as well as C^*) when choosing between helping relatives or helping strangers.

The conditional obligations approach can be spelled out either: 1. in terms of obligations that are conditional on helping, or 2. in terms of obligations that are conditional on incurring a cost (the idea here is that one can justify doing A rather than C because of the cost involved in C , yet one cannot appeal to this cost to justify doing B rather than C , given that both options involve this cost). Both interpretations run into difficulties. Since all four options involve helping others, it is of no use to bring in obligations that are conditional on helping. Similarly, since the agent is not incurring any costs in this situation, bringing in obligations that are conditional on incurring costs is to no avail.¹⁷

6 Agent-sacrificing options

The cost incurred by x in doing B or C has to be sufficiently large relative to the benefit conferred upon y for it to be permissible that x does nothing (= A).¹⁸ From the agent-relative perspective, the cost to x has to be at least as large as the benefit

¹⁶Unlike accounts that are based on a self v. other asymmetry, e.g. Sider: 1993 and Portmore: 2008, and accounts that build in a separate principle of partiality, e.g. Hurka and Shubert: 2012 and Horton: 2017.

¹⁷In addition, conditional accounts lack the resources to address the dynamic inconsistencies identified in section 7, since this requires a suitable set of orderings in terms of which strategies can be classified as being suboptimal by a global evaluation that takes counterfactual alternatives into consideration.

¹⁸What classifies as sufficiently large depends on the size of the benefit. There is no absolute threshold but a relative threshold. If the benefit is very large, then an agent can be required to incur a large cost. The agent is required to incur the cost as long as $\beta > k \times \alpha$.

to y : $V_{\text{relative}}(A) \geq V_{\text{relative}}(C)$, i.e. $k \times \alpha \geq k \times 2 \times \beta$.¹⁹ Otherwise, A will not be a top-element in the agent-relative ordering. The cost is also not allowed to be too large (relative to the benefit), i.e. $\alpha \leq \beta$. The benefit would otherwise not be worth the cost from the agent-neutral perspective. Otherwise, A would be the (only) top-element in the agent-neutral ordering and B would be impermissible when choosing between A and B.

Accordingly, the following inequalities have to be satisfied:

$$k \times \alpha \geq 2 \times \beta > \beta \geq \alpha$$

(When using the weak dominance condition, A needs to be the only top-element in the agent-relative ordering if it is to be permissible. The agent-relative evaluation then needs to satisfy a slightly stronger condition: $V_{\text{relative}}(A) > V_{\text{relative}}(C)$ and hence $k \times \alpha > 2 \times \beta$. Likewise, A is not allowed to be a top-element in the agent-neutral ordering for B to be permissible when choosing between A and B. The agent-neutral evaluation then needs to satisfy a slightly stronger condition: $V_{\text{neutral}}(B) > V_{\text{neutral}}(A)$ and hence $\alpha < \beta$.)

Since $\alpha \leq \beta$, agent-relative prerogatives only allow for departures from what is impartially best by justifying agent-favouring options whereby the agent can be partial to himself and those who are close to him. They do not, however, justify agent-sacrificing options whereby the cost that is incurred by the agent is greater than the benefit that is conferred.²⁰ Such behaviour is suboptimal from both perspectives. It can neither be condoned from the agent-neutral perspective nor from the agent-relative perspective.

The prohibition on agent-sacrificing behaviour can be explained by noting that the well-being of the agent matters morally²¹ and generates reasons in the same way as the well-being of anyone else.²² The benefits hence have to be sufficiently large in order to justify the costs that the agent has to incur. The costs to the agent speak against the action and there have to be sufficiently strong counter-vailing considerations that speak in favour of incurring these costs. Cases of agent-sacrificing whereby the benefit is smaller than the cost do not satisfy this condition. Engaging in agent-sacrificing is not merely imprudent but morally problematic. By bringing about less value, both agent-relatively and agent-neutrally,

¹⁹Given that y is a stranger, the agent-relative modifier for y is: $k = 1$. We hence have the weak inequality: $k \times \alpha \geq 2 \times \beta$.

²⁰Thanks to Tom Hurka for pressing me to say more about agent-sacrificing options.

²¹Agent-relative good is not a form of prudential good but of moral good. The agent-relative v. agent-neutral distinction is not to be confused with the prudential v. moral contrast. What is good relative to x is also not to be conflated with what is good for x . This is particularly clear when concerned with benefitting people who are close to the agent. Benefiting them is not a matter of prudence (or some other form of non-moral reason) but a matter of partiality.

²²This account differs sharply from theories based on a self v. other asymmetry. Portmore, for instance, argues that “there is no requiring reason to promote one’s own wellbeing per se” (Portmore: 2008. p. 416) but instead only a justifying reason.

than if the action had not been performed, such an agent is not suitably responsive to moral reasons.

Difficulties would arise in making sense of agent-relative prerogatives if one were to treat the agent's well-being differently. If one identifies the agent-neutral point of view as the moral point of view, then it is difficult to see how it can be permissible to depart from this point of view.²³ Accordingly, we should not think of prerogatives in terms of non-moral reasons competing against moral reasons. Instead, we are dealing with two moral points of view from which things can be evaluated. Since neither perspective is privileged, these two points of view must agree that an action is dominated by another alternative for that action to be impermissible.

One can respect that one's own happiness gives rise to moral reasons yet nevertheless incorporate agent-sacrificing options into the theory by including a patient-relative ordering. This ordering corresponds to an evaluation that is centred on the patient – one person's agent-relative ordering is another person's patient-relative ordering. One can then apply dominance principles over the triple.

WEAK TRIPLE DOMINANCE

$$C(X) = \{x \in X : \neg \exists y \in X [(y \geq_{\text{agent}} x \wedge y \geq_{\text{neutral}} x \wedge y >_{\text{patient}} x) \vee (y \geq_{\text{agent}} x \wedge y >_{\text{neutral}} x \wedge y \geq_{\text{patient}} x) \vee (y >_{\text{agent}} x \wedge y \geq_{\text{neutral}} x \wedge y \geq_{\text{patient}} x)]\}$$

STRICT TRIPLE DOMINANCE

$$C(X) = \{x \in X : \neg \exists y \in X (y >_{\text{agent}} x \wedge y >_{\text{neutral}} x \wedge y >_{\text{patient}} x)\}$$

This makes it permissible not only to act on the basis of the agent-relative evaluation alongside the agent-neutral evaluation, but also to occupy the position of another. The agent is allowed to put himself into the shoes of someone else, namely the person that the agent is acting upon, and act on the basis of the evaluation centred on that person's perspective. If that person could permissibly bring about a certain distribution of costs and benefits by maximising that person's agent-relative ordering, then it is likewise permissible for the agent to do so, i.e. to maximise the patient-relative ordering.²⁴ (Given a suitable set of agent-relative constraints, one can explain why it is permissible for the agent to sacrifice himself in order to benefit someone else, yet not permissible for the stranger to impose the same costs on the agent in order to incur the same benefit.)

²³This problem is closely related to the 'paradox' of supererogation. Cf. "It is hard to see how it could be permissible, from the moral point of view, to refrain from doing something that you have an undefeated reason (from that very point of view) to do. ... Supererogation, according to this way of seeing things, turns out to be impossible" (Dreier 2004, 148).

²⁴The dominance principles rule out suboptimal agent-sacrificing (both the gratuitous conferring of a smaller benefit and the gratuitous incurring of a greater cost). Such behaviour is classified as suboptimal on the patient-relative ordering.

This allows for limited agent-sacrificing (contrary to the self-other asymmetry). Although the costs that the agent can permissibly incur can be greater than the benefits that are thereby conferred, the costs are nevertheless not allowed to be excessive. Instead of $\alpha \leq \beta$, the condition is now that $k_{z,x} \times \alpha \leq k_{z,z} \times \beta$.²⁵ The fact that the modifier k applies to α ensures that how much x can sacrifice in order to benefit z depends on how close z is to x . How the cost α is evaluated from the patient's perspective depends on the distance between z and x , which in turn depends on the nature and strength of the morally significant relationships in which z stands to x . The closer z is to x , the more significant α is from z 's perspective. Relative to a given benefit β that accrues to z , the cost α that can be incurred by x will be smaller the closer z is to x . Those who are close to the agent cannot as easily condone self-sacrificial behaviour since the agent's well-being matters more from their perspective. In the same way that benefits to those who are close to the agent matter more from the agent-relative perspective, so costs to the agent matter more from the patient-relative perspective when the beneficiary is close to the agent.²⁶ More is at stake from z 's perspective so that the sacrifice is more easily classified not only as agent-relatively as well as agent-neutrally suboptimal but also as patient-relatively suboptimal.

Difficulties arise if multiple people are affected by the agent's actions. Once multiple patients are involved it becomes unclear whose perspective the agent can take up. Can the agent identify with particular individuals (or particular pluralities of individuals)? Or is the only patient-perspective that the agent can permissibly adopt that of all the patients taken collectively? If the former were the case, then the prohibition of suboptimal beneficence would be significantly weakened.

	x	y
A	o	o
B	$-\alpha$	$+\beta$
C	$-\alpha$	$+(2 \times \beta)$

In the original case, B would still be impermissible. However, the situation would be different in a variant case in which the beneficiary of the suboptimal act of beneficence differs from the beneficiary of the optimal act of beneficence.

	x	y	z
A'	o	o	o
B'	$-\alpha$	$+\beta$	o
C'	$-\alpha$	o	$+(2 \times \beta)$

²⁵The first k corresponds to the distance between the patient and the agent: $k = f(d(z, x))$, whereas the second is $\max k = f(d(z, z))$. Since d does not satisfy the symmetry axiom, it is a pseudo-metric. Morally significant relationships need not be symmetric: x can be relatively close to z , without z being close to x .

²⁶If z is required to incur a cost $-\alpha$ to confer a benefit $+\beta$ on x , where $\alpha < \beta$, then x cannot sacrifice $-\beta$ in order to confer $+\alpha$ on z (assuming the distance between x and z is symmetrical).

The agent's ranking is: $A' > C' > B'$ and the neutral ranking is: $C' > B' > A'$. Even though the ranking centred on the patients taken collectively is likewise: $C' > B' > A'$, the relative ranking centred on the individual patient y is: $B' > C' > A'$. When taking this final ranking into consideration, all three alternatives are permissible. Unlike B in the original case, B' will then not be impermissible. This suggests that the relevant contrast is between the agent (the one doing the acting) and the patients taken collectively (all those who are acted upon). What the agent can do is to take up their collective perspective rather than identify with particular individuals.

One then has to find a way of combining their perspectives and construct a plural perspective of the various patients taken collectively. We can do this by allowing the distance function to take plural arguments, letting the multiplier k be a function of the average distance. The distance of the zz 's to the agent x , i.e. $d(zz, x)$, which determines the modifier k that applies when evaluating x 's well-being from the perspective of the zz 's will be the average distance:

$$\frac{\sum_{z_i \prec zz} d(z_i, x)}{|zz|}$$

7 Dynamic inconsistency

A and B are permissible when $X = \{A, B\}$, yet only A and C but not B are permissible when $X' = \{A, B, C\}$. Expansion consistency condition BETA, accordingly, fails. This condition requires that if two options are in a given choice set, then one of them is in the choice set of any superset of the original set of alternatives iff the other is also in the choice set.

$$\text{BETA} : [x, y \in C(X) \wedge X \subseteq X'] \rightarrow [x \in C(X') \leftrightarrow y \in C(X')]$$

Whilst violating condition BETA, contraction consistency condition ALPHA is satisfied.²⁷

$$\text{ALPHA} : [x \in C(X) \wedge x \in X' \subseteq X] \rightarrow x \in C(X')$$

This violation of BETA is analogous to that which arises when dealing with incommensurable values, where two non-comparable options can be in a choice set yet only one of them fails to be in the choice set of a superset of the set of alternatives, since this superset involves a further alternative that dominates only this option but is non-comparable to the other. Although the agent-relative axiology does not involve incommensurability, since the agent-relative and agent-neutral evaluations are connected by a modification function and hence are comparable, we are nevertheless dealing with choice functions that are defined over two separate orderings.

²⁷Expansion consistency conditions GAMMA and DELTA are likewise satisfied.

The fact that agent-relative prerogatives imply BETA-violations is disconcerting since this gives rise to dynamic inconsistencies.²⁸ If an agent makes use of myopic choice principles (which are purely forward-looking) or sophisticated choice principles (that use backwards induction), then there are cases in which he can, by means of a sequence of permissible choices, bring about an outcome that is deemed to be impermissible from the outset. The prohibition on suboptimal beneficence thus appears to conflict with dynamic consistency.

This already came up in the sequential choice situation in which an agent is faced on two occasions with the possibility of either doing nothing or conferring benefit β at cost α . A dynamic inconsistency when applying the pure (rather than mixed) dominance principle.²⁹ Only the pure strategies, namely 1. not helping in both cases or 2. helping in both cases, are permissible from the outset. However, when evaluated at the second choice node, both helping and not helping are permissible independently of how the agent has chosen at the first choice node (as long as one is using choice principles that satisfy separability). Only pure strategies are admissible at the outset, yet mixed strategies can permissibly result when applying the choice principles afresh at the second choice node. This dynamic inconsistency is avoided by means of the mixed dominance principles, since they classify mixed strategies as being permissible at the outset.

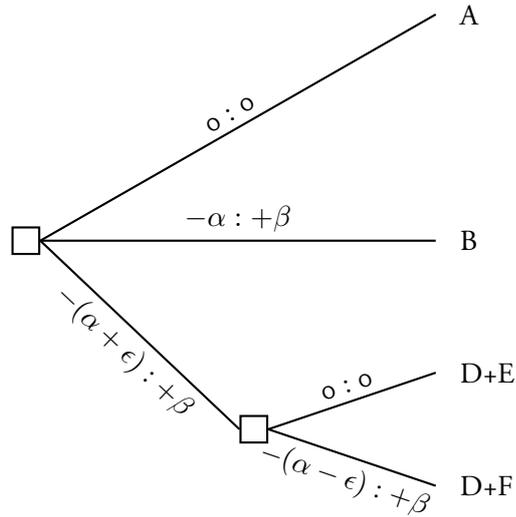
Some problems, however, remain. Dynamic inconsistency looms large in the following choice situation. Agent x has to choose between A , B and a third option $D =$ save y 's left arm at a cost slightly greater than α , but then face a subsequent choice between $E =$ do nothing and $F =$ save y 's right arm at a cost slightly less than α , i.e. $D+E$ is the dynamic equivalent of B^- , whereas $D+F$ is the dynamic equivalent of C .³⁰

²⁸BETA-violations also have implications for conditional obligations. When $X = \{A, B, C\}$ the choice set is $C(X) = \{A, C\}$. One is required to do either A or C . Given that one is not going to do C , it might seem that one is required to do A , i.e. that one is subject to a conditional obligation to do A on condition that one is not going to do C . This would be counterintuitive, cf. Horton: 2017, p. 94. However, given that BETA is violated, this conditional obligation does not follow. The choice set of a restriction of the set of alternatives need not correspond to the (non-empty) restriction of the choice set. This is brought out clearly by an alternative formulation of BETA that highlights the corresponding contraction consistency condition: $X \subseteq X' \wedge C(X) \cap C(X') \neq \emptyset \rightarrow C(X) \subseteq C(X')$. Instead of restricting the choice set, i.e. $C(X|¬C) = C(X) \setminus C = \{A\}$, we have to apply the choice function to the restricted set of alternatives, i.e. $C(X|¬C) = C(X \setminus C) = \{A, B\}$. We then end up with the conditional obligation to do either A or B on condition that one is not going to do C . (If the evaluation fails to be context-free, then the fact that C is a possible option can affect the evaluation of B such that $C(X|¬C)$ need not be the same as $C(X \setminus C)$.)

²⁹Sider's self-other utilitarianism (cf. Sider: 1993), for instance, suffers from this kind of dynamic inconsistency.

³⁰The same reasoning applies if one alters the benefit conferred on y , rather than the cost to the agent:

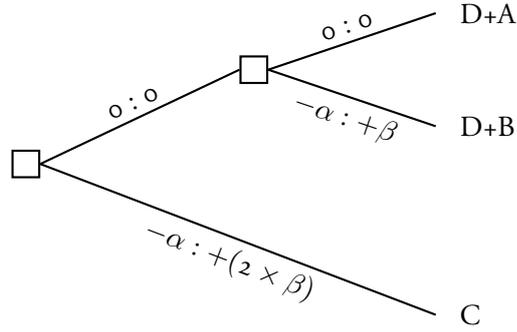
	x	y
A	o	o
B	$-\alpha$	$+\beta$
D+E	$-(\alpha + \epsilon), o$	$+\beta, o$
D+F	$-(\alpha + \epsilon), -(\alpha - \epsilon)$	$+\beta, +\beta$



D+F is impartially best. D+E, by contrast, is suboptimal both agent-relatively and agent-neutrally. Since it is dominated on both orderings by B, it is impermissible according to the mixed dominance principles. Hence, one is only permitted to choose D rather B at n_0 given that one is going to do F rather than E. Yet, the agent can permissibly perform D+E when choosing myopically, i.e. when applying the choice principle afresh after having chosen D and when faced with the choice between E and F. Once the agent has done D and is choosing between the continuations E and F, the additional cost that was involved in D is a sunk cost that does not affect this subsequent choice. The evaluation at the original node n_0 does not cohere with that at node n_1 . We are thus confronted with dynamic inconsistency. The strategy D+E is impermissible at the outset, yet can be implemented by means of a sequence of permissible choices.

A case similar to that considered in section 2 is even clearer. This time the agent does not first decide whether to incur the cost but instead decides whether to do C or do nothing. If the agent decides to do nothing at n_0 , he then faces a subsequent choice at n_1 of again doing nothing or choosing B.

	x	y
A	o	o
B	$-\alpha$	$+\beta$
D'+E'	$-\alpha, o$	$+(\beta - \epsilon), o$
D'+F'	$-\alpha, -\alpha$	$+(\beta - \epsilon), +(\beta + \epsilon)$



The strategy D+B is impermissible from the outset. Yet it can be implemented by means of a sequence of permissible choices. In the isolated tree both A and B are permissible. Separability then implies that B is a permissible continuation of D at node n_1 , conflicting with the evaluation at n_0 . The separability principle, accordingly, has to be restricted if one is to ensure dynamic consistency. It has to be weakened so that the choice between A and B in the isolated tree can conflict with the choice in the subtree.

Fortunately, we can restrict separability such that it is compatible with the global choice principle developed in the next section, yet strong enough to underwrite the argument in section 2 against suboptimal beneficence. Restricted separability applies if the maximal elements (i.e. those that are undominated with respect to the intersection quasi-ordering) are comparable in terms of the intersection quasi-ordering, i.e. the agent-relative and agent-neutral orderings agree on how they are to be ranked (which happens if they stand in \equiv to one another). If a subtree is such that the maximal elements at n_i are comparable with each other, then the choice in the subtree has to agree with the choice at n_0 in a corresponding isolated tree. The members of the choice set are comparable in that case and the maximising and optimising choice functions yield the same choice set. This does not presuppose that all alternatives are comparable (i.e. it is not required that $\forall x, y \in X(x \gg y \vee x \equiv y \vee x \ll y)$).³¹ This is because it is possible for two non-comparable alternatives to be both dominated by the same option, such that any x, y in X are either comparable or both dominated by a further option that is comparable to each of them. The maximal elements can thus be comparable to all other elements, without the non-maximal elements having to be comparable with respect to each other.³²

³¹This only follows if comparability is an equivalence relation.

³²In fact, we can accept a slightly stronger condition: any maximal subset of comparable members of the choice set at node n_i will be the choice set at node n_0 of an isolated tree consisting of all those options that are in the set of alternatives X and that are comparable to the options in the subset of the choice set of X .

8 Global choice

To ensure dynamic consistency one has to appeal to global choice principles that consider the entire decision tree and not only those nodes that are accessible from the current choice node.³³ Global choice principles are required because the choices are not independent. Unlike in the situation where the agent is faced on two occasions with the possibility of helping and where the choices are independent, in the case of {A, B, D+E, D+F} the choices are connected. The additional cost involved in choosing D rather than B is incurred precisely so that one can choose a certain option later on. The rationale for choosing D is to be able to choose F later on. There is no point in choosing D unless one subsequently chooses F. Choosing E directly contravenes this rationale. By undermining the rationale for D, choosing E makes the prior choice irrational and pointless.

Choosing E at n_1 is impermissible because doing so makes having chosen D impermissible. D is only permissible as part of the D+F combination but impermissible as part of the D+E combination. D inherits its (im)permissibility from the (im)permissibility of the course of action of which it is a part. Accordingly, whether one chooses E or F determines whether or not D is impermissible. The principle that it is impermissible to do something that renders a past action impermissible then implies that, even though E and F are permissible when considered by themselves, only F is permissible qua continuation of D.

Unlike in many other cases of dynamic inconsistency, sophisticated choice is not the solution. A sophisticated chooser cannot avoid inconsistencies arising from incompleteness by using backward induction in order to anticipate what he is likely to choose in future choice situations and then decide amongst available alternatives in light of the anticipated choice behaviour. After all, both E and F are permissible when evaluated myopically. There is thus no reason to think that the agent will choose one rather than the other. Instead, the choice function has to take past choices into consideration. One has to choose in light of what choices one has already made and what options were previously available.

Dynamic inconsistency arises when the evaluations at n_i and n_j conflict, i.e. when there is a difference between what was deemed to be permissible and what is now deemed to be permissible (where this is neither due to a change in information nor a change of objectives nor a difference in terms of means-ends reasoning, but entirely due to a change in perspective). This kind of conflict can be avoided by switching from local to global evaluations. As long as deliberation and evaluation are sensitive to the past, the evaluations at the different nodes will cohere. A

³³This should be unsurprising given that incompleteness in general requires global choice. E.g. myopic choice principles allow one to zigzag through a sequence of non-comparable alternatives to end up with drastically suboptimal outcomes, even when dealing with a local form of incompleteness that is to be understood in terms of an intersection quasi-ordering generated from two dimensions of evaluation that can be integrated in different ways. Cf. "Incompleteness and dynamic consistency" (Bader: manuscript).

change of temporal perspective has no effect on an all-encompassing evaluation that is not only future-directed but also takes the past into consideration on an equal footing. The different evaluations will, accordingly, not conflict.

The thought underlying global choice is that agents are diachronic beings. We should be concerned with the agent considered as a whole. We should evaluate choices not only locally, but also globally by asking how they fit into the agent's life as a whole. By basing one's choice on the most encompassing evaluation, one ensures that one is maximally responsive to reasons, not qua synchronic time-slice, nor qua forward-looking agent, but qua unified diachronic agent. Privileging the diachronic agent over the present time-slice as well as over the future segment of the agent explains why the past can normatively constrain the present evaluation.

Although an agent can only choose amongst alternatives that are currently available, he can evaluate these alternatives by comparing them to other alternatives that were previously available. Comparison, unlike choice, need not be restricted to available alternatives. It can also encompass counterfactual alternatives. Previously available options can cast a historical shadow. One has to expand the comparison class by considering counterfactual alternatives alongside available alternatives when deciding which of the latter to choose. An evaluation that considers options that the agent could have chosen is not restricted to local features of the choice situation but considers the choice in light of the agent's life as a whole. This diachronic perspective allows alternatives that are undominated by the available alternatives to nevertheless be suboptimal due to being dominated by options that the agent could have chosen.

Counterfactual options matter because what the agent does in the present can affect our evaluation of the agent's past choices. We do not only evaluate isolated actions but also complex courses of actions. An action ϕ can be permissible as part of one course of action, yet impermissible as part of another. How we are to evaluate x 's having ϕ -ed can thus depend on what the agent does at later times. The agent's subsequent actions determine whether ϕ -ing is extended in such a way that a permissible course of action results. This kind of impermissibility only shows up from the diachronic perspective. Past choices matter from this perspective, rather than being sunk costs that are to be ignored. The past matters for diachronic evaluations. Although the past is fixed, i.e. the agent cannot change the fact that he has ϕ -ed, the evaluation of ϕ -ing is not fixed. For instance, the agent can affect its evaluative status by extending it in a way that renders having ϕ -ed impermissible.

When including counterfactual alternatives in the evaluation, an available alternative can be dominated by something on a different branch of the tree that is no longer available but that was accessible at a prior choice node. One considers not only available alternatives that one can choose, but also counterfactual alternatives that one could have chosen. In the case at hand, E is not dominated by any of the available alternatives (although F is agent-neutrally better, it is agent-

relatively worse). However, when E is considered as a continuation of D, it is dominated by the counterfactual alternative B. Although E is not dominated by an available alternative, it is dominated by an alternative that could have been performed had one chosen differently at a prior choice node. As a result, F is the only undominated available alternative from the perspective of global choice and hence is the option that is to be chosen. Dynamic inconsistency is thus avoided.

The set X_i contains the alternatives that are accessible from n_i . Relative to a subsequent node n_j some of these alternatives will be merely counterfactual, namely those in $X_i \setminus X_j$. They are not in X_j , but could have been reached had the agent acted differently at n_i . Although these alternatives are no longer accessible and hence cannot be chosen, they can nevertheless be used for evaluating available alternatives. Being undominated in X_j , i.e. $\neg\exists y(y \in X_j \wedge y \gg x)$, is then a necessary but not sufficient condition for ending up in the global choice set. In addition, alternatives have to be undominated in a suitably expanded set X_i that also contains counterfactual alternatives.

The choice set relative to a given counterfactual expansion is selected by means of a myopic choice function $C_M(X^{j-n})$. It contains those members of X_j that are undominated relative to the counterfactual expansion resulting from going back n steps in the decision tree. We select alternatives from X_j on the basis of being undominated by those in X_i where $j - i = n$.

$$C_M(X^{j-n}) = \{x : x \in X_j \wedge \neg\exists y(y \in X_i \wedge y \gg x)\} \text{ where } j - i = n$$

We go back in the decision tree and expand the set of options with which the available alternatives are compared. Available alternatives are compared to counterfactual alternatives that could have been reached from the choice node that precedes X_j by n steps. The further back we go, the more counterfactual alternatives are included. Since contraction consistency condition ALPHA is satisfied, it follows that what is dominated in one set will also be dominated in every superset.

$$C_M(X^{j-n}) \subseteq C_M(X^{j-m}) \text{ where } n > m$$

The global choice function evaluates whether available alternatives are undominated in a suitably expanded set containing counterfactual alternatives (where this includes the degenerate expansion in which no alternatives are added). The global choice function C_G , accordingly, does not select an option that is not likewise selected by the myopic choice function C_M . The global choice function considers a superset of alternatives, yet can only select available but not counterfactual alternatives from this superset. Since the myopic choice function C_M selects precisely from these available alternatives, the choice sets stand in a subset relation.

$$C_G(X) \subseteq C_M(X)$$

There cannot be a switch from an alternative being unchosen when choosing myopically to being chosen when applying the global choice function. Instead,

it is only possible for something to be chosen by a myopic agent yet rejected by a global agent. Accordingly, if the myopic choice set is a singleton, then there is no need to consider any counterfactual alternatives. Any superset X^{j-n} which includes counterfactual alternatives and for which the choice function $C_M(X^{j-n})$ yields a nonempty choice set will contain this unique alternative.

$$C_M(X) = \{x\} \rightarrow C_G(X) = \{x\}$$

If some alternative is undominated in X° , which contains the alternatives that could have been reached from the original choice node, then the global choice set contains all options that are undominated in this maximal counterfactual expansion.

$$C_M(X^\circ) \neq \emptyset \rightarrow C_G(X) = C_M(X^\circ)$$

Problems arise when the agent has made a mistake and picked a dominated alternative at some point in the decision tree. Mistakes ensure that $C_M(X^\circ) = \emptyset$. Every available alternative is then dominated by some counterfactual alternative. This can happen in two ways. Either an agent makes a myopic mistake by choosing an option that is dominated by available alternatives. One goes down a branch that is such that all the outcomes that can be reached from that branch are dominated by outcomes that can be reached by taking a different branch. Or an agent makes a global mistake by choosing an option that, though not being dominated by available alternatives, is nevertheless dominated by counterfactual alternatives (whereby the agent did not make any myopic mistakes at the choice nodes involved in the counterfactual expansion).

If the agent has made a mistake and there are a number of different undominated alternatives such that $C_M(X_j)$ is not a singleton, then we cannot go back to the original choice node. Instead, we can only go back to the choice node n_{j-m} that has a non-empty choice set, i.e. $C(X^{j-m}) \neq \emptyset$ and that is such that all preceding nodes, i.e. X^{j-n} where $n > m$, yield an empty choice set. In the case of a myopic mistake at node n_i one can only expand up to node n_{i+1} . By contrast, in the case of a global mistake at node n_i one can expand up to and including node n_i but not beyond that. We then choose myopically from X^{j-m} .³⁴

$$C_G(X_j) = C_M(X^{j-m}) \text{ such that } C_M(X^{j-m}) \neq \emptyset \wedge C_M(X^{j-n}) = \emptyset \text{ for all } n > m$$

The options that remain after the set of counterfactual alternatives has been maximally expanded, without rendering the choice set empty, are the actions that are to be chosen and that are maximally rational from a diachronic perspective.³⁵

³⁴ $C_M(X^{j-n})$ where $n > j$ is defined as \emptyset , since the relevant set of counterfactual alternatives is undefined. This ensures that one goes back to the original choice node, such that the choice set is $C_M(X^\circ)$, if the agent has not made any mistakes by the time he reaches node n_j .

³⁵Thanks to Roger Crisp and especially to Kacper Kowalczyk for helpful comments. Thanks also to participants of the Arizona Workshop in Normative Ethics and the Oxford Population Ethics seminar.

References

- [1] BADER, R. M. Kantian Axiology and the Dualism of Practical Reason. In *The Oxford Handbook of Value Theory*, I. Hirose and J. Olson, Eds. Oxford University Press, 2015.
- [2] BADER, R. M. Conditions, modifiers, and holism. In *Weighing Reasons*, E. Lord and B. Maguire, Eds. Oxford University Press, 2016, pp. 27–55.
- [3] CUBITT, R. Rational dynamic choice and expected utility theory. *Oxford Economics Papers* 48, 1 (1996), 1–19.
- [4] DREIER, J. Why ethical satisficing makes sense and rational satisficing doesn't. In *Satisficing and Maximizing – Moral Theorists on Practical Reason*, M. Byron, Ed. Cambridge University Press, 2004, pp. 131–154.
- [5] HORTON, J. The all or nothing problem. *The Journal of Philosophy* 114, 2 (2017), 94–104.
- [6] HURKA, T., AND SHUBERT, E. Permissions to do less than the best: a moving band. *Oxford Studies in Normative Ethics* 2 (2012), 1–27.
- [7] KAGAN, S. *The Limits of Morality*. Oxford University Press, 1989.
- [8] MCMAHAN, J. Doing good and doing the best. In *Philanthropy and Philosophy: Putting Theory Into Practice*, P. Woodruff, Ed. Oxford University Press, 2017.
- [9] PARFIT, D. Future generations: further problems. *Philosophy and Public Affairs* 11, 2 (1982), 113–172.
- [10] PORTMORE, D. Dual-ranking act-consequentialism. *Philosophical Studies* 138 (2008), 409–427.
- [11] PUMMER, T. Whether and where to give. *Philosophy and Public Affairs* 44, 1 (2016), 77–95.
- [12] SEN, A. Maximization and the act of choice. *Econometrica* 65, 4 (1997), 745–779.
- [13] SIDER, T. Asymmetry and self-sacrifice. *Philosophical Studies* 70, 2 (1993), 117–132.