The asymmetry

Ralf M. Bader

ABSTRACT: This paper provides an account of the asymmetry in population ethics. The first half of the asymmetry is explicated by means of a person-affecting view, whereas the second half is established by means of a structural consistency constraint. This account can be integrated into a general theory that can handle (i) cases where there are externalities in that members of the original distribution are positively or negatively affected by bringing the miserable life into existence, (ii) cases in which one is concerned not only with bringing individual persons into existence but also groups of people, and (iii) situations in which it is uncertain whether an action will result in the addition of lives that are worth not living.

1 The basic account

The asymmetry in population ethics:

- The fact that a life would be worth living does not, by itself, generate reasons to bring it into existence.
- The fact that a life would be worth not living does, by itself, generate reasons to not bring it into existence.[1,2]

Accounting for this intuition is one of the key desiderata of any satisfactory theory of population ethics. The asymmetry, however, is difficult to explicate. Attempts at making sense of it run into a dilemma. On the one hand, a person-affecting view can account for the first half, on the basis that it is not better for the happy person to exist than not to exist. Yet it struggles with the second half. Given that a miserable life is not worse than non-existence, a person-affecting view has difficulty in explaining what could speak against bringing such a life into existence. On the other hand, an impersonalist view can deal with the second half. Miserable lives should not be brought into existence because their existence is taken to

[1] The qualification 'by itself' implies that we set aside effects on others in evaluating reasons for adding/not adding lives.

[2] Terminological note: lives that are worth living are also referred to as 'happy lives' and lives that are worth not living as 'miserable lives'.

be impersonally bad, i.e. bad from the point of view of the universe rather than that of any individual. However, it is unable to explain the first half, given that the impersonal goodness of lives that are worth living likewise generates reasons. As a result, the asymmetry is difficult to sustain.

We have two unpalatable symmetrical options that either over-generate or under-generate reasons. Either existence and non-existence are comparable with respect to some value, in which case one has reason not to create miserable lives but then also has reason to create happy lives, leaving one with an impersonalist axiology that violates the neutrality intuition. Or they are non-comparable, in which case one does not have reason to create happy lives but then ends up not having reasons not to create miserable lives, leaving one with a radical person-affecting view (such as that espoused by Heyd: 1992) that has counter-intuitive commitments (cf. McMahan: 2009).

	IMPERSONAL	PERSON-AFFECTING	ASYMMETRY
worth living	R(add)	¬R(add)	¬R(add)
worth not living	R(¬add)	¬R(¬add)	R(¬add)

Impersonalist theories consider the existence of happy lives to make things better, in the same way that the existence of miserable lives makes things worse. Yet, if adding happy lives makes the world a better place, then one has reason to add such lives. This means that in order to underwrite the first half of the asymmetry, impersonalist approaches would have to find a way of making these reasons disappear.[3] They would have to explain why the value of happy lives does not provide reasons, whereas the disvalue of miserable lives does provide reasons. Making value count only in the one case but not in the other is difficult and conflicts with plausible bridge-principles connecting values and reasons.

Person-affecting approaches, by contrast, do not have to make any reasons disappear. The problem that they have to address is to explain why there are reasons not to add miserable lives, even though there are no axiologically-based reasons that speak against adding miserable lives. This is much less problematic, since one can bring in additional resources that supplement axiologically-based reasons in order to explain the asymmetry. Person-affecting theories thus have a crucial advantage over impersonalist theories when it comes to underwriting the asymmetry, since one can explain why there are reasons against ϕ-ing in a situation in which there are no axiologically-based reasons against ϕ-ing, by appealing to supplementary non-axiologically-based reasons, whereas one cannot explain why there are no reasons in favour of ϕ-ing in a situation in which there are axiologically-based reasons in favour of ϕ-ing.

[3] The first half of the asymmetry states not merely that it is permissible not to add such lives (which could be explained in terms of prerogatives), but that there is no reason, not even a pro tanto reason, to add them.

1.1 The first half

The first half of the asymmetry follows from the intuition of neutrality. According to this intuition, adding a person neither makes a distribution any better nor any worse but is instead axiologically neutral.[4]

This intuition can be explained by means of an account of conditional goodness, according to which the goodness and reason-giving force of a person's well-being is conditional on the existence of the person. We can compare the goodness of existing lives and order them in terms of how good they are. Yet, we cannot compare existence with non-existence. We cannot compare something that has goodness with something that lacks goodness. Given that goodness is conditional on existence, we should be neutral about existence. Existence is not better than non-existence. Nor is non-existence better than existence. Yet, existence and non-existence are not equally good either. Instead, existence and non-existence are not comparable. No betterness relations hold between existence and non-existence. Existence is not better than non-existence and there is hence no reason to bring a happy life into existence. Instead of having reason to bring lives into existence, one only has reason to improve lives that exist.[5] "We are in favor of making people happy, but neutral about making happy people" (Narveson: 1973, p. 80).

The idea of conditional goodness and the attendant commitment to the non-comparability of situations in which the condition is satisfied with those in which it fails to be satisfied applies straightforwardly to considerations of personal good. According to non-comparativism, it is not possible to compare existence with non-existence from the point of view of personal good. Any attempt to make a comparison to the effect that it is better for a person to exist than not to exist is confused and misguided. The personal betterness relation only holds between situations in which the person in question exists. One situation can be better for x than another situation only if x exists in both situations. Accordingly, existence is not better for x than non-existence. These scenarios are not comparable with respect to x's good.[6]

If x exists, then it is better for x to be happy than not to be happy. Yet, it is not

[4] Adding a person can have positive or negative effects on other people, yet such effects are set aside when examining whether the existence of the person makes things better or worse by itself.

[5] The notion of existence is not to be understood in a temporal manner, i.e. that we have reason to improve a life after it has come into existence. Instead, it is to be construed atemporally. We have reason to improve lives if their existence is given and not contingent on our actions (i.e. every member of the set of alternatives is such that it contains the lives in question).

[6] The personal betterness relation is a dyadic relation that has lives as its relata. In particular, the 'better for x'-relation is a dyadic relation that compares x's lives. When comparing existence with non-existence, one of the two relata is missing, which ensures that the relation cannot apply. Moreover, there is a lack of good-making features in the case of non-existence. Non-comparativism thus follows from the fact that there cannot be a betterness relation without relata and from the fact that there cannot be goodness without good-making features. Cf. "The neutrality of existence" (Bader: manuscript).

better for x to exist and be happy than for x not to exist. That x's life goes well is not better for x than that it is not lived at all. Non-existence is neither better nor worse for x than either a happy or a miserable existence. The difference between happiness and misery makes no difference when these states are compared with non-existence, since both of them fail to be comparable with non-existence. If the condition is not satisfied, then goodness is not applicable. The value of a non-existent life is undefined. It lacks value rather than having zero value. The field of the personal betterness relation is, accordingly, not complete but restricted to existence cases. Situations falling within the field are not comparable to ones falling outside it. The normative significance of happiness/misery, i.e. the ability of these features to make a difference, is thus conditional on existence.

The rejection of comparability in terms of personal good leaves open whether one can compare existence with non-existence from some other point of view, such as the point of view of general good or impersonal good. If one were to adopt an impersonal approach, for instance, then one could hold that, even though it does not make sense to ask whether x's existence is better for x than x's non-existence, it is intelligible to inquire as to whether x's existence is impersonally better than x's non-existence. Although the state of affairs in which x exists is not comparable in terms of personal good, it may well be comparable in terms of general or impersonal good. A person-affecting approach, however, forecloses this possibility. If a person-affecting view is adopted, then non-comparativism regarding the personal betterness relation leads to unrestricted non-comparativism.

A person-affecting approach requires all betterness claims to ultimately be reducible to claims about how persons are affected. Distributions are to be evaluated in terms of how they affect various people and the goodness of distributions is reducible to the goodness for the members of the distributions.[7] The reducibility requirement is stronger than a mere supervenience or functional dependence requirement. The fact that D_1 is better than D_2 has to consist in facts about personal betterness. Put differently, the betterness relation between distributions has to consist in nothing other than the obtaining of personal betterness relations amongst the members of the distributions. Facts about general good have to be reducible to facts about personal good, such that facts about general good just are plural facts about personal good. The general betterness relation is then to be construed as a plural comparative.[8]

[7] What is known as the 'person-affecting restriction', namely that a distribution D_1 cannot be better than another distribution D_2 unless there is someone for whom D_1 is better than D_2, is a necessary condition on betterness orderings that is both too strong and too weak. It is too strong, since it mistakenly builds in a rejection of impartiality. It is too weak, since it can be satisfied by impersonal theories that merely assign ethical relevance but not ethical significance to personal good. The positive commitment of a person-affecting view, as it is understood in this paper, is the commitment to the reducibility of general good and the rejection of impersonal good.

[8] For a detailed development and defence of this person-affecting approach, cf. "Person-affecting population ethics" (Bader: manuscript). A crucial upshot of this approach is that distri-

Two key commitments follow from this reducibility requirement, namely (i) that considerations of general good are countenanced only as long as they are reducible to considerations of personal good, and (ii) that considerations of impersonal good are rejected altogether on the grounds of being irreducible. All claims of goodness are then either directly about the personal goodness of individual lives or reducible to such claims. Non-comparativism about personal good together with the reducibility of general good implies non-comparativism with respect to general good. Together with the rejection of an impersonal standpoint, this implies that one ends up with there not being any point of view from which existence can be compared with non-existence. In short, existence cannot at all be compared with non-existence because (i) non-comparativism about personal good ensures that such a comparison is not possible from the personal point of view, (ii) the reducibility of general good rules out such comparisons from the general point of view, and (iii) the rejection of impersonal good excludes the possibility of a comparison from the impersonal point of view.

Adopting a person-affecting approach thus ensures that one ends up with unrestricted non-comparativism. Conditional goodness, consequently, applies both to personal good and to general good. According to both the personal and the general betterness relation, it is better that a person who exists is happy rather than not happy. Given that someone does exist, that person's happiness matters from the point of view of personal good as well as from the point of view of general good. Yet, it is not better according to either of these betterness relations that a happy person exists than that such a person does not exist.

This commitment to conditional goodness explains the first half of the asymmetry.[9] Since the goodness of a person's happiness is conditional upon the existence of the person, it follows that the addition of a happy person does not make things better. Instead, it has a neutral impact, where neutrality is to be understood in terms of non-comparability, given that we cannot compare a situation in which the condition is satisfied with one in which it fails to be satisfied. The fact that a life would be worth living, consequently, does not by itself generate reasons for bringing it into existence.

1.2 The second half

Accounting for the second half of the asymmetry is not so straightforward. The commitment to neutrality is unrestricted. The addition of a life is axiologically

butions are only comparable when they are equinumerous.

[9] In addition, one has to argue that non-comparative considerations do not generate reasons to bring happy lives into existence, e.g. that people can be non-comparatively benefited by being brought into existence since, even though a life that is worth living is not better than non-existence, it is nevertheless good for the person (cf. Parfit: 1981, Appendix G; McMahan: 2013). For an argument that non-comparative considerations cannot speak in favour of bringing happy lives into existence cf. "The neutrality of existence" (Bader: manuscript).

neutral in the case of all possible lives and not only in the case of lives falling within a certain range. Existence is non-comparable with non-existence, no matter how good or bad the life in question is. In the same way that it is not better that a happy life is lived, it is not worse that a miserable life is lived. As a result, there seems to be nothing that speaks against bringing miserable lives into existence.

To avoid this unpalatable situation, one has to appeal to non-axiological resources.[10] This section explains why we have reasons not to bring miserable lives into existence in terms of a structural consistency constraint that actions have to satisfy. This allows us to combine a symmetrical account at the axiological level with an asymmetry at the level of reasons, without undermining standard bridge-principles connecting values and reasons and without having to bring in independent considerations, such as justice-, fairness-, or rights-based considerations.[11]

To begin with, we need to consider what it is for a life to be worth not living.[12] Given non-comparativism, one cannot explicate this notion in terms of a life being worse than non-existence. One cannot simply compare a life with non-existence and evaluate whether it is better for the life to be lived than not to be lived, i.e. whether it is better for the person to exist than to not exist, such that a life is worth living if it is better than non-existence, neutral if it is equally as good as non-existence, and worth not living if it is worse than non-existence. Instead of a life worth not living being one that is worse than non-existence, it is one that is equally as good as some member of the following set:

$$\{L : \forall t \exists t' [t' < t \wedge V(L_{o-t'}) > V(L_{o-t})]\}$$

Every life in this set is such that for any point in time t in that life the value of the life up to that point, i.e. $V(L_{o-t})$, is smaller than the value of the life up to some earlier point t'. We cannot say of these lives that it would have been better

[10] Bringing in non-axiological resources is unproblematic since the asymmetry concerns the level of reasons, not the level of values.

[11] Substantive non-axiological considerations are problematic. They do not straightforwardly apply to non-existent entities: for instance, it is dubious to hold that there can be a right against being brought into existence since there cannot be a rights-bearer that can have this right: as long as the person does not exist there is no rights-bearer and hence no right, yet once the person does exist, that person cannot have a right not to be brought into existence, since it is impossible to make it the case that something that already exists never comes into existence. Moreover, they are likely to yield reasons against creating miserable lives that are too strong and that cannot be adequately integrated into a general theory. One needs to be able to weigh up the reasons not to add miserable lives against other reasons in order to deal with (i) cases involving externalities whereby other people are positively or negatively affected by the addition of the life, (ii) cases in which groups of people are added, some of whose lives are worth living whilst others are worth not living, and (iii) cases in which there is uncertainty as to whether a particular action will result in the addition of lives that are worth not living (cf. sections 2.2–2.4).

[12] The notion of a life worth living is an axiological notion based on the notion of well-being. It is concerned with what is good/bad for the person and is to be distinguished from the notion of a meaningful life. That a life is not worth living does not imply that it is meaningless.

if they had never been lived, but only that it would have been better had they ended sooner rather than later: the shorter they are, the better they are.

If we partition all lives into equivalence classes under the relation of being equally as good as and use the betterness relation to impose an ordering on these equivalence classes, then the class of neutral lives is the least upper bound of the set of classes that have the members of the aforementioned set as members, whilst those lives that are better than neutral lives are lives that are worth living.

In this way, it is possible to explain what it is for a life to be worth not living in terms of the internal structure of the life, in particular in terms of the axiological ordering of various initial segments of the life. This account does not compare existence with non-existence but only compares different shortenings of the life. One compares shortenings of the life with other shortenings, without having to appeal to the non-sensical idea that it is worse for the life to be lived than for the life not to have been lived at all.[13]

A paradigm case of a miserable life is one that is constantly getting worse, in the sense that lifetime well-being is constantly decreasing. For every point in this life, the value of the life up to that point is smaller than the value of the life up to each earlier point. Consider such a life L that occupies a temporal region that is topologically open on the left. Every shortening of such a miserable life is likewise a miserable life. Moreover, for every shortening s, there exists a prior shortening s* such that $s^* > s$. When we start with a distribution D_1 and consider adding L, then, although the original distribution D_1 is not comparable with the extended distribution $D_2 = D_1 \cup L$ since they are not equinumerous, D_2 is comparable with distributions $D_{2'}$, $D_{2''}$ ... that contain various shortenings of L, e.g. $D_{2'} = D_1 \cup s'$. These extensions of D_1 form an open-ended series of distributions that has the original distribution as its limit.

For instance, if we consider a sequence of shortenings in which each member of the sequence is half as long as the previous member, i.e. the length of the n^{th} shortening is given by $\frac{1}{2^n}$ times the length of the added life, then the limit of this

[13] This account does not have atomistic presuppositions and does not make assumptions about separability across time, but is instead compatible with a holistic understanding of the value of a life. The underlying idea is that the goodness of a life depends on what happens in the life in such a way that the shorter the life is the less room there is for value in this life. This means that though there can be holistic goods, such as pattern goods, these goods have to depend on the length of the life in such a way that they matter less the shorter the life is. In short, every good and bad (whether atomistic or holistic) has to diminish with duration, such that an infinitesimal life can only contain an infinitesimal amount of value. Accordingly, there can be no duration-invariant holistic goods, i.e. holistic goods that are equally good no matter how long the life lasts and hence do not diminish or disappear as the life gets shorter and shorter. Such goods would not show up under shortening and would not generate reasons against shortening.

(If separability were accepted, one could identify a neutral time-slice as one that is such that the life with this time-slice is equally as good as the life without it. A neutral life can then be understood as one that is equally as good as a life that only consists of neutral time-slices.)

sequence is:

$$\lim_{n\to\infty} \frac{1}{2^n} = 0.$$

As n tends to infinity, the life being added gets shorter and shorter, until in the limit it no longer exists. The original distribution is not a member of this sequence of ever shorter extensions and is not better than any member of this sequence. Instead, it is the limit of the sequence. The sequence of shortenings is such that, in the limit, the length of the added life is zero, i.e. $\frac{L}{\omega} = 0$, which means that no life is added and that D_{2^ω} is identical to D_1. In the limit, the life is not lived at all and the distribution is not extended. In other words, the original distribution is the limit of the sequence of extensions resulting from adding ever shorter shortenings of L.[14]

Such a paradigmatically miserable life is worse than all its shortenings. Moreover, each shortening is worse than all prior shortenings. The life gets better and better the shorter it is. Correspondingly, the extended distribution gets better and better the shorter the life is that is being added. This means that someone who can choose, not only to either add or not add this life, but also has the option to add various shortenings of the life, will choose the option of not adding the life. If one enriches the set of alternatives by including all the possible shortenings into the set of alternatives, i.e. X = {L, $\frac{L}{2}$, $\frac{L}{4}$, …, ¬add}, then the only admissible option is to not add the miserable life, i.e. C(X) = {¬add}. This holds not only for lives that are constantly getting worse but applies to all lives that are such that for every shortening s, there exists a prior shortening s* such that $s^* > s$. In all these cases the life is made better and better by being made shorter and shorter, such that, in the limit, it is not lived at all.

An idealised agent who has the ability to shorten the life at any moment in time, i.e. can stop the effects of the action from unfolding, and who is ideally responsive to moral reasons cannot add a miserable life.[15] The commitment to shorten the life whenever doing so makes it better subsumes an infinite sequence consisting of ω-many actions that jointly prevent the life from coming into existence.

If t is the boundary on the left of the temporal region of the life, then a commitment to shorten the life whenever doing so makes it better implies that the life neither exists at t nor at any time after t. On the one hand, time t is the boundary of the life and is not part of the temporal region occupied by the life.

[14] The closeness measure is a temporal rather than an evaluative measure. Closeness is understood in terms of the length of the life, not the value of the life. We are dealing with a sequence of initial segments that has as its temporal limit the non-existence of the life, such that the corresponding sequence of actions, namely the sequence of earlier and earlier shortenings, likewise has as its causal limit the non-existence of the life.

[15] In the rest of this section miserable lives will be understood as lives that are never worth living. Lives that, though worth not living overall, nevertheless start out being worth living will be considered in section 2.1.

On the other, if the life were to exist at $t + \frac{1}{2^n}$, then there would already have been a prior member of the sequence of shortenings, say number $n + 1$, which would have cut the life short at $t + \frac{1}{2^{n+1}}$. Put differently, were the life to exist at any time t'' after t, there would have been an earlier time t' (where $t < t' < t''$) at which the life would already have been cut short, contradicting the claim that the life exists at the later time t''. Since every time after t is such that there is a prior time at which the life would already have been cut short, it follows that the life cannot exist at any time after t. Given that adding the life implies that it exists at some time after t, and given that the life cannot exist at any such time, it follows that adding a miserable life is logically precluded by the commitment (joined with the ability) to shorten it whenever one has reason to do so.[16]

There is thus a contradiction between 1. a miserable life is added and hence exists at some time after t, 2. for any time t'' in the life there is an earlier time t' such that it is better for the life to be shortened at t' than to continue until t'', and 3. for every time t' if it is better to shorten the life at t' than to let it continue, then it is shortened at that time. Adding a miserable life is, accordingly, incompatible with an effective commitment to performing at every point in time all those shortenings that would make the life better and that one would hence have reason to perform, i.e. $\forall t[\text{R(shorten at } t) \rightarrow \text{shorten at } t]$. This general commitment subsumes ω-many conditional commitments (shorten at t_1 if the life exists at t_1, shorten at t_2 if the life exists at t_2 ...). Adding a life is admissible only if these commitments do not preclude the possibility of adding the life, i.e. $\text{add} \in C(X)$ only if $\Diamond(\text{add} \wedge \forall t[\text{R(shorten at } t) \rightarrow \text{shorten at } t])$.

The phenomenon that we are interested in differs importantly from cases in which the choice-set is empty when using a maximising or optimising function

[16] If the temporal region occupied by the life were topologically closed on the left (rather than open), then the sequence of shortenings would prevent the life from existing at any time after t but would not prevent it from existing at t. Instead of being prevented from coming into existence, there would instead be either a point-sized life occupying an extensionless temporal region, or a temporally extended life that occupies a temporal region corresponding to the minimum extent of a life and hence is such that none of its initial segments would classify as lives. For the argument to work, lives have to occupy regions that are topologically open on the left, since the sequence of shortenings would otherwise be closed. This happens naturally if time (or space-time) is not pointy but gunky, in which case every region has sub-regions, making it the case that there are no temporal points but only nested sequences of ever smaller extended temporal regions. Additionally, this can happen if lives cannot be instantaneous but must have positive extent without there being any minimum extent. If there is no minimal unit having axiological significance, then this implies that every axiologically significant unit is temporally extended and has temporally extended axiologically significant units as parts, i.e. the structure of axiologically significant units will be gunky if these conditions are satisfied, which can be the case even if time (or space-time) should fail to be gunky. (Furthermore, one can argue that even if there could be point-sized lives or lives that occupy a temporal region that is closed on the right and that are such that none of their initial segments classify as lives, these lives could not be miserable lives on the grounds that misery has to be temporally extended, without there being any minimal extent. This would still be sufficient for establishing the fine-grained asymmetry discussed in section 2.1.)

due to there not being any maximal element in the betterness ordering.[17] If there were an open-ended sequence on the positive side, for instance if there were an infinite number of possible extensions of a happy life, then there would not be any maximal element in the sequence of extensions. Yet no incompatibility would be generated.

The key difference between the two scenarios is that in the former case the limit of the betterness ordering coincides with the limit of the temporal/causal sequence, since always extending amounts to adding either an infinitely long life or a life of finite duration (if the sequence asymptotically approaches a limit). In the latter case, by contrast, the evaluative limit comes apart from the limit of the temporal/causal sequence. The limit of the betterness ordering consists in a neutral life. The limit of the temporal/causal sequence, however, is not a neutral life. In fact, it is not a life at all. In the limit, no life is added and the distribution is not extended, since always shortening amounts to not adding the miserable life at all. Put differently, whilst $\mathrm{L} \times \omega$ is simply a long life, $\frac{\mathrm{L}}{\omega}$ is no life at all and is accordingly equivalent to $\neg$add. The problem is thus not that there is no maximal element, in that for every shortening there is a yet better shortening, but that the shortenings collectively preclude there being a life at all.

This difference can be brought out clearly if the idealised agent can perform supertasks, such that the set of alternatives includes $\mathrm{L} \times \omega$ in the former situation and $\frac{\mathrm{L}}{\omega}$, which equals $\neg$add, in the latter situation. Hence, the choice set will be $\mathrm{C}(\mathrm{X}) = \{\neg\text{add}, \mathrm{L} \times \omega\}$ when dealing with a happy life that can always be further extended and thereby be made better and better, but only $\mathrm{C}(\mathrm{X}) = \{\neg\text{add}\}$ when dealing with a miserable life, thereby rendering the addition of the miserable life inadmissible. Whereas always extending a life is compatible with there being a life, namely one that is infinitely long or that asymptotically approaches a limit, always shortening a life is incompatible with the existence of a life since, in the limit, no life is lived, thereby rendering only the latter but not the former incompatible with the addition of a life.

When we evaluate actions from an axiological perspective, we usually evaluate their effects and compare the outcome of one action with that of another. We can, however, also evaluate the differences between outcomes, instead of evaluating the outcomes separately. The differences that ϕ-ing makes vis-à-vis ψ-ing can be evaluated and compared with various shortenings of those effects, where a shortening is a partial evaluation, i.e. one only evaluates the effects up to some time t. In the case of adding a life, the alternative is not to add the life. The

[17] It also differs from cases in which performing all the actions in a sequence of ω-many actions leads to sub-optimal outcomes, despite the fact that each action in the series is such that one has reason to perform that action (cf. Arntzenius & McCarthy: 1997, pp. 49-50 and Barrett & Arntzenius: 1999), requiring one to adopt a satisficing rather than maximising or optimising strategy. As Scott and Scott: 2005 have shown, such infinite exchange problems arise only given certain conditions involving the individuation and tracking of the units that are exchanged, none of which apply in the context that we are presently considering.

difference that the action makes is thus the existence of the life and the various shortenings of the effects of the action are nothing other than shortenings of the life. One then evaluates not only the whole life but also various initial segments of the life. Importantly, the shortenings of the differences that ϕ-ing makes vis-à-vis ψ-ing can be comparable, even when the outcomes of ϕ-ing and ψ-ing fail to be comparable.

An action ϕ is inadmissible if it violates the shortening consistency constraint. This constraint is violated if there is some alternative ψ such that the difference that ϕ-ing makes vis-à-vis ψ-ing is such that there is a sequence of shortenings of these effects that starts with the degenerate shortening (which consists in the complete eventuation of the effects that result from ϕ-ing rather than ψ-ing) and that has the complete non-occurrence of these effects as its limit, i.e. the outcome of ψ-ing is the limit of the sequence of shortenings of the effects of ϕ-ing, whereby for every shortening s there exists a prior shortening s* that is better than s.

The underlying idea is that one should not perform an action that is such that one should prevent its effects from eventuating (if one were able to do so). Whilst ϕ-ing is compatible with preventing some of its effects from unfolding, it is incompatible with preventing all of its effects from unfolding. In the case at hand, adding a life is compatible with cutting it short at some later point, yet adding a life is incompatible with cutting it short at every point (of a sequence that has the non-occurrence of the action as its limit). If an action is such that its effects should be completely precluded from unfolding, then this action should not be initiated.[18]

Adding a miserable life violates the shortening consistency constraint. There is a sequence of shortenings, each of which makes the life better than the subsequent shortenings, that has the non-existence of the life as its limit. The different shortenings jointly preclude the possibility of adding a miserable life. The shortening consistency constraint, accordingly, requires one to refrain from adding such a life, even though doing so is not worse than the alternative of not adding the life. This means that the reason against adding miserable lives is explained in terms of the very feature that makes them worth not living, namely that they are better the sooner they end. It is the internal axiological structure of a life that both explains why the life is worth not living and that leads to the violation of the consistency constraint and thereby makes it inadmissible to add such a life.

[18] If the outcomes of ϕ and ψ are comparable, then the shortening consistency constraint rules out ϕ-ing whenever its outcome is dominated by that of ψ-ing (given the coarse-grained approach of section 2.1). The consistency constraint thus rules out all dominated alternatives. This means that the shortening consistency constraint is not an ad hoc principle that is brought in specifically to deal with the asymmetry, but is a general constraint that generates a choice function that selects all maximal elements that conform to the asymmetry. Every theory needs a choice function that gets us from the level of values/reasons to oughts. The shortening consistency constraint constitutes precisely such a criterion: it generates a same-number maximisation requirement together with the asymmetry in one fell swoop.

By contrast, no analogous reasoning can be used in the case of lives that are worth living to show that one has reason to bring them into existence. If we add lives that are worth living, then all that can happen is that it is better for such lives to continue rather than be cut short. Yet the fact that one should prevent such lives from being shortened does not generate a reason to bring them into existence in the first place. Accordingly, if the life is worth living, then it is both admissible to not extend the distribution and admissible to add the happy life. Both actions are admissible. They do not violate any consistency constraints and there are no alternative actions that are better than them. Given their non-comparability, one should be neutral between these options. The choice whether or not to extend the distribution by adding the happy life is then an existential choice.

Whereas the fact that adding a miserable life violates the shortening consistency constraint renders that action inadmissible and gives us reason not to perform it, the fact that adding a happy life satisfies this consistency constraint does not give us any reason to perform that action, but merely makes it the case that it is an admissible action. This is what explains the asymmetry.

2 Extending the account

2.1 Lives that are initially worth living

The basic account applies to any miserable life that is such that for every point in that life t, there exists a prior point t' that is such that the life up to t' is better than the life up to t. Such lives only constitute a sub-class of all the lives worth not living, namely those that are never worth living (though they may be worth continuing at various points). Those lives that, considered as a whole, are worth not living even though they are worth living or neutral up to some point do not fall within the scope of the argument. In order to end up with an unrestricted asymmetry, the account needs to be extended to apply, not only to those lives that are such that every point in the life is such that it would have been better had the life been shorter, but also to lives that are worth living up to some point yet worth not living on the whole.

A life is no longer worth living at time t if $\exists t'(t' \leq t \wedge \forall t''[t' < t'' \rightarrow V(L_{o-t'}) > V(L_{o-t''})])$.[19] The point after which a life is no longer worth living is the first point that is such that the life up to that point is better than the life up to all later points, i.e. the point at which it reaches its peak in terms of lifetime well-being.[20] A life that is worth living up to some point yet that is worth not living

[19] A life, considered as a whole, can be worth living even though it reaches a point at which it is no longer worth living. The badness to be found in the life after that point simply has to be sufficiently insignificant as to not outweigh the goodness that has accrued up to that point.

[20] If the life has multiple peaks or a flat peak, then it is the final point at which it peaks that is such that after this point the life is no longer worth living, since only this final point is such that

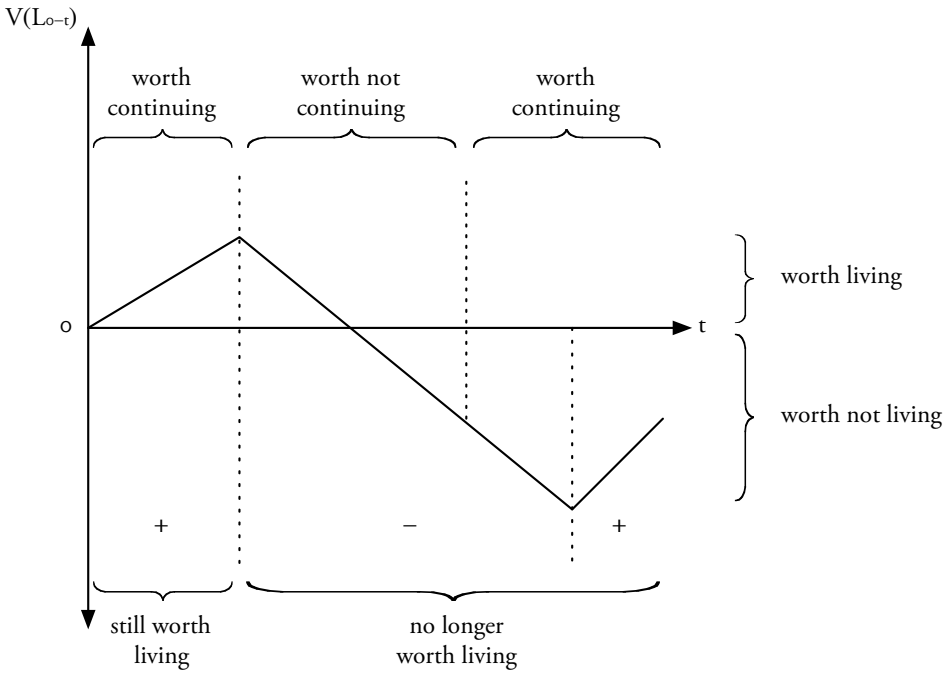

considered as a whole is thus a miserable life that has the property of not being such as to be no longer worth living at all points in time. This means that there are points in time at which it is still worth living, yet such that the (net) goodness of any initial segment is not sufficiently large to outweigh the (net) badness accruing to the life during later segments when it is no longer worth living.

coarse-grained options

Given that a life that is worth not living is such that any goodness that accrues whilst it is still worth living will be outweighed by the badness accruing later on, it follows that there exists a sequence of shortenings of the life that has as its limit a situation in which the life is not lived at all, where each of the shortenings is one that makes the life better than the life would be if it were to continue to the end. Each shortening is such that the life lived up to that time is better than the life lived until the end.

The shortening consistency constraint thus only needs to be slightly weakened to preclude the addition of any miserable life (and not only of a restricted range of miserable lives). Rather than assessing whether there is a sequence of shortenings such that for every shortening s there exists a prior shortening s* that is better than s, one assesses whether every shortening s is better than ϕ, such that there is reason to shorten the life rather than let it continue to the end. Since ψ-ing rather

the life is better up to that point than up to all later points, whereas prior peaks fail to satisfy this condition due to there being later peaks that are equally good.

than ϕ-ing amounts to preventing the differences that ϕ-ing makes vis-à-vis ψ-ing from eventuating, one should ψ rather than ϕ, when the differences that ϕ-ing makes are to be cut short at every point of a sequence that has the complete non-occurrence of these effects as its limit rather than be allowed to fully eventuate. Accordingly, one should not add rather than add miserable lives, even when these lives are initially worth living.

Adding a miserable life implies that it exists at some time after t (where t is the boundary on the left of the temporal region of the life). However, for every time t'' after t, there exists a prior time t' (where $t < t' < t''$), at which the life would already have been cut short, since the life up to t' is better than the life lived to the end, thereby ensuring that it cannot exist at t''. Since it is precluded from existing at any time t'' after t, the addition of the life is precluded. Appealing to coarse-grained options, where one can either shorten the life or let it continue until the end but where one cannot decide to let the life continue until some later time at which it stops being worth living, allows one to extend the asymmetry to miserable lives that are worth living up to some point. If we compare a sequence of choices involving coarse-grained options, in which one does not compare different shortenings with each other, but rather compares the different shortenings with the option of letting the life continue to the end, then each choice-situation in the sequence will be such that it will be better to shorten the life than to let it run to the end.[21] Adding a miserable life thus violates the weakened shortening consistency constraint, since one will be required to shorten the life rather than let it continue to the end in each choice-situation, which in the limit amounts to not adding the life.

The weakened consistency constraint ensures that the asymmetry applies to all miserable lives, including those that are worth living up to some point. All miserable lives are such that there is a sequence of shortenings that has the original distribution as its limit and that is such that it is always better to shorten the life than to let it continue to the end. By contrast, any such sequence in the case of happy lives would have as its limit the point after which the life first becomes better than the life lived to the end and hence does not violate the shortening consistency constraint. Accordingly, no miserable lives are to be added, whereas one is to be neutral about adding happy lives.

FINE-GRAINED ASYMMETRY

Alternatively, one can adopt a fine-grained construal of the asymmetry, according to which the fact that a life is worth not living only provides us with reasons not to extend it beyond the point after which it is no longer worth living but not reasons to not add the life. One should be neutral about adding such lives, but not about extending them beyond the point after which they are no longer worth

[21] Since these different choice situations cannot be faced successively by one agent, they have to be construed counterfactually.

living. The fact that a life would be miserable does not suffice to speak against adding it, since a miserable life can be worth living up to some point. It is only miserable lives that are never worth living that one has reason not to add at all, since in the case of those lives every point is such that the life is worth not living.

On this construal the asymmetry is not really about miserable lives but about lives that fail to be worth continuing after some point. The limit of the relevant sequence of shortenings will be the point t in the life after which it is no longer worth living. Every point t'' after that point will be such that there is a prior point t' (where $t < t' < t''$) that is such that the life up to t' is better than the life up to t''. What matters for the argument is that there is a point after which things go badly (in the sense that any continuation of the life after that point will have a negative net impact), independently of whether the misery accruing after this point is sufficient to outweigh the initial happiness and thereby make the life miserable on the whole, or whether it merely detracts from the value of the life without making it the case that the life is worth not living. In other words, the issue will be whether the life is still worth living, not whether the life is worth living on the whole. The difference between happy and miserable lives only becomes pronounced in the case of those miserable lives that are never worth living since such lives are not to be added at all, whereas this can never happen when it comes to happy lives since all such lives are such that one should be neutral about adding them (though not about extending them beyond the point after which they are no longer worth living).

2.2 Externalities

Discussions of the asymmetry usually hold everything fixed except for the addition of a miserable life. Correspondingly, the principles meant to explain the asymmetry are generally qualified by ceteris paribus conditions. Yet, it is not clear how these principles are to be extended to situations in which other things are not equal, for instance situations in which members of the original distribution are positively or negatively affected.

Dealing with these more complicated situations is difficult for theories that bring in deontological resources. They have to treat these deontic considerations as being commensurable with the relevant axiological considerations, so that they can be combined to yield an overall verdict. Appealing to substantive non-axiological resources thus generates problems and complicates the task of integrating the asymmetry into an overall normative theory. The shortening consistency account, by contrast, can be integrated into an overall theory that deals with situations in which ceteris is not paribus.

Prima facie, the proposed account might seem to consider the reason not to add miserable lives as being absolute, given that this reason is established, not on the basis of the disvalue of the resulting state of affairs, but on the basis of a violation of a consistency constraint. As such, it would seem to classify as a con-

straint rather than as a pro tanto consideration that could be outweighed. Since the reason not to create miserable lives is established by appealing to a consistency constraint, it is unclear what determines the strength of this reason and how it can vary with the degree of miserableness. It is hence unclear how one is to integrate the considerations deriving from the asymmetry with other considerations to end up with an overall verdict that specifies whether one has reason not to add the miserable life.

While the degree of miserableness is not required to establish the asymmetry (since the reasoning equally applies to lives that are only slightly miserable as it does to lives that are extremely miserable), the extent of the miserableness of the lives to be added matters when considering the effects that adding these lives has on other people. In particular, one can consider whether the reason for shortening is sufficiently strong to outweigh any positive externalities that might result from the addition of the miserable life. The question then is whether the good that is done by shortening the miserable life is sufficient for outweighing the positive effects that accrue to others, i.e. whether the amount of suffering prevented by shortening the life is greater than the benefits that accrue to others. The miserableness of the created life thus plays a role in determining how much happiness accruing to others can be outweighed to ensure that shortening is not only better as far as the miserable life is concerned but better all-things-considered.[22,23] When considering not only how much the shortenings improve the miserable life but also taking externalities into consideration, it becomes possible for the reason in favour of shortening to be outweighed. If the benefits are sufficiently significant, then there will not be a violation of the shortening consistency constraint from the perspective of an all-things-considered evaluation. In this way, the reason against adding miserable lives turns out not to be absolute.

If adding a miserable life brings about sufficient benefits to others, then it can turn out that we do not have reason against bringing it into existence. Accordingly, we should be neutral about adding miserable lives with sufficiently large positive externalities. Correspondingly, if adding a happy life brings about sufficient suffering to others, then it can turn out that we do have reason not to bring such a life into existence. This means that we should not be neutral about adding happy lives with large negative externalities. We should not create happy people if the negative effects on others are sufficiently significant.[24] However, it

[22] In this context it is important to distinguish between axiologically generated reasons and deontic considerations regarding compensation.

[23] The added happiness to others does not have to result directly from extending the distribution but can be an indirect result or be due to a common cause. (It can even be entirely disconnected, though in such cases it is natural to suppose that the agent faces a further alternative in which this additional disconnected happiness is to be found but which does not include any added lives.)

[24] Given that effects on other people are taken into consideration, this claim neither conflicts with the asymmetry nor with the intuition of neutrality, since they are both restricted claims that abstract from effects on others and consider the added life in isolation.

will never be the case that we have reason to extend a distribution because of the additional happiness that results for members of the original distribution. Such happiness only provides reasons when considered from the point of view of those who are benefitted, but not from the moral point of view that takes everyone into consideration.

2.3 Adding groups

When considering adding groups of people, one needs to find a way for the happiness of happy lives to weigh against the misery of miserable lives. Otherwise, one miserable life would be sufficient for making the addition of any group of people morally problematic, no matter how many happy lives this group should contain. The happiness of those whose lives will be worth living must be able to neutralise the reason not to add the group deriving from the miserable life (cf. Sikora: 1978, p. 137).

The shortening consistency account can achieve this result by assessing whether it would be better if the miserable life were shortened, even if there were to be a corresponding shortening of the other lives being added. If there is a sequence of combined shortenings of the different lives that has the original distribution as its limit and that is such that each shortening makes the lives considered collectively better, as happens for instance when one is faced with the option of adding twins one of whom will be highly miserable while the other will only be moderately happy, then the shortening consistency condition implies that one should not bring this group into existence.[25] By contrast, if the happy lives being added contain sufficient happiness to outweigh the misery to be found in the miserable lives, then there will not be such a sequence of shortenings.

2.4 Probabilistic cases

In the context of certainty, a unique outcome is associated with each alternative. When dealing with uncertainty, a number of possible outcomes can result from a particular action. The question then is what happens when there is a likelihood that ϕ-ing will result in the addition of a life that is worth not living. Under what conditions are there reasons against performing actions that have a chance of bringing about miserable lives? What is needed, in particular, is an account that can explain how the possibility of a life being happy can cancel out or weigh against the possibility of a life being miserable, in order to avoid the unpalatable commitment that any non-zero probability of a life being miserable would be sufficient to generate (overall) reasons against bringing it into existence.

[25] If the combined shortenings are defined in relative terms, e.g. where shortening s_n involves each of the added lives $L_1, L_2 \ldots$ being shortened to $\frac{L_i}{2^n}$, then all the comparisons amongst the members of the sequence of shortenings will be same-number (in fact same-people) comparisons.

Whilst cases of uncertainty pose difficulties for deontic asymmetries invoking rights- or justice-based considerations, given that we lack a clear account of the conditions under which the imposition of a risk amounts to a rights-violation, the shortening consistency account does not run into difficulties. The explanation of the asymmetry applies equally when working with expected value.

When acting under uncertainty, the question is whether shortening the life is expected to make things better such that one has reason to commit to shorten the life. This can be established by comparing the expected outcome with shortening to the expected outcome without shortening.[26] One has to evaluate whether for every time t, there is a prior time t′, such that the expected value of the life up to t′ is greater than the expected value up to t, i.e. whether $EV(L_{o-t'}) > EV(L_{o-t})$. In this way, one can apply the consistency principle to the expected value of shortenings of the life. In the case of uncertainty, the consistency constraint has to be understood, not in terms of shortenings that will make the life better, but instead in terms of shortenings that are expected to make the life better (and to which an agent who is fully responsive to reasons would make an ex ante commitment). Rather than assessing the shortenings that would make the life better, one assesses the shortenings that one has reason to commit oneself to performing to see whether they jointly preclude the addition of the life.

There are two ways of aggregating probabilistically discounted evaluations of the different possibilities to determine the expected value of shortenings of a life.

1. TIME-SLICE AGGREGATION

First one aggregates across states of nature and then across time. One determines the expected temporal well-being at different times and then aggregates these to arrive at the expected lifetime well-being (up to time t).

$$EV(L_{o-t}) = \sum_{o<j\leq t} EV(L_j) = \sum_{o<j\leq t} \sum_{i\in N} p_i \times V(L^i_j)$$

Time-slice aggregation runs into difficulties.

1. This way of aggregating presupposes the separability of times and hence precludes holistic evaluations of the value of lives.

   It might be suggested that one can apply a holistic value function to the expected time-slice profile, such that $EV(L_{o-t})$ is not determined by simply summing the expected time-slice values between o and t, but instead is determined by $f[EV(L_i)$ for all $i \in o < i \leq t]$, where f is some non-separable holistic aggregation function. This, however, generates wrong results since the holistic function is to be applied to the well-being profiles in particular

[26] The different shortenings can be characterised either in terms of absolute units, i.e. shorten after n units, or in terms of proportions relative to the longest life, i.e. shorten at $\frac{L}{n}$.

states of nature. It is not the expected shape of the life that matters but the various possible shapes that the life can have in the various states of nature. For instance, if evenness matters, then (other things being equal) one should prefer a lottery over even distributions to a lottery over distributions that are uneven but that are such that the unevenness is cancelled out in the process of aggregation so that the expected shape of the life is even. This happens for instance in the case of a lottery with a 50% chance of a constantly increasing life and a 50% chance of a mirror-reversed constantly decreasing life. In this kind of case, one knows that the latter lottery will result in an uneven distribution and that one will consequently not get evenness, despite the expected time-slice profile being even.

2. This account cannot handle lives of different lengths. It leads to non-comparability since L_j will not be defined if L does not exist at time j. This, in turn, precludes aggregation, insofar as $\sum p_i \times V(L^i_j)$ will not be defined if not all of the states of nature contain a life that exists at j.

2. INITIAL-SEGMENT AGGREGATION

One aggregates first across time and then across states of nature. One determines lifetime well-being (up to t) in the various states of nature and then aggregates their probabilistically discounted values.

$$EV(L_{o-t}) = \sum_{i \in N} p_i \times V(L^i_{o-t})$$

The expected value of the life up to t is given by the aggregate of the probabilistically discounted values of the life in the various states of nature.[27] This account does not presuppose the separability of times. The value function applied to initial segments of the life can fail to be separable across times. Nor does it run into difficulties when the lives of the added person are of different lengths in different states of nature. Lifetime well-being up to t, i.e. L_{o-t}, is defined even if the life ceases to exist at some prior time $t' < t$. Accordingly, we need to aggregate probabilistically discounted lifetime well-being, rather than aggregate expected time-slice well-being.

DIFFERENT-NUMBER CASES

The account can straightforwardly deal with same-number cases, where one is adding a life (or a set of lives) such that it is a matter of chance whether it (or they) will be worth not living.[28] Difficulties arise in different-number cases, when it is a matter of chance how many (if any) people will come into existence whose lives

[27] This account assumes that states-of-nature are separable. Whilst the independence axiom is plausible, one can build in a non-separable aggregation function over states of nature, i.e. $EV(L_{o-t}) = f[p_1 V(L^1_{o-t}), \ldots, p_n V(L^n_{o-t})]$.

[28] Given a commitment to impartiality, all that matters is the number of lives that are being

will be worth not living, i.e. the different possibilities associated with a particular action involve different numbers of people.

Probabilistic cases involving different numbers seem to require aggregating non-equinumerous distributions. This, however, cannot be done by means of balancing (the aggregative procedure compatible with person-affecting views) since there will not be a bijective mapping with respect to which gains and losses are defined. Accordingly, the expected value of the distribution under various shortenings will not be defined for different-number cases and one cannot assess whether shortening is expected to make things better.

Although we cannot aggregate all the possible distributions associated with a particular risky action, we can supervaluate over them. If every possible distribution resulting from an action is such that it should not be brought about, then the action should not be performed. Moreover, we can partition the possible distributions into equivalence classes and then aggregate probabilistically discounted equinumerous distributions. If every equivalence class of possible distributions resulting from an action is such that it should not be brought about, then the action should not be performed.

Mixed cases, where some classes are neutral and others generate reasons against extending the distribution, are more difficult. The question then is whether some equivalence class being such that there is reason against it being brought about is sufficient for ruling out that action. Neither possibility is satisfactory. Either there being some equivalence class suffices for ruling out the action, in which case any non-zero chance of there being miserable lives that are non-comparable with the other possible outcomes will render an action impermissible, thereby over-generating reasons against extending the distribution. Or it does not suffice and instead all equivalence classes have to satisfy this condition, in which case any chance, no matter how small, of there being a number of happy lives that is non-comparable with the other possible outcomes will render an action permissible, thereby under-generating reasons against extending the distribution.

The question that needs to be addressed is whether one has reason to commit to shorten at t if there is a chance that it will make things better and a chance that it will make things worse, whereby the distributions with respect to which these betterness judgements are made cannot be aggregated. To answer this question one has to determine how strong the reason is to shorten, i.e. by how much shortening is expected to make things better in the states of nature in which a distribution of size m will result, and how strong the reason is not to shorten, i.e. by how much shortening is expected to make things worse in the states of nature resulting in a distribution of size n. Whether one should commit to shorten

added and not their identities. This means that the account can deal with cases, in which it is certain that one person will be brought into existence as a result of a particular action, but where it is uncertain whether it will be x or y that will be added, since one can abstract from the identity of the person.

depends on the comparative strength of these reasons.

For instance, an idealised agent would make an ex ante commitment to shorten at t' in a mixed case involving both m-sized distributions where shortening is expected to make things better and n-sized distributions where it is expected to make things worse if $EV(D^m_{o-t'}) - EV(D^m_{o-t}) > |EV(D^n_{o-t'}) - EV(D^n_{o-t})|$. This comparison only involves value-differences. Even though m-sized distributions are not comparable to n-sized distributions, differences between m-sized distributions can be compared to differences between n-sized distributions. Such comparisons do not require different-number comparability. There can be unit-comparability without level-comparability of distributions from different equivalence classes.[29] Even though it does not make sense to ask in mixed cases whether the distribution is expected to be better as a result of shortening, one can ask whether the ex ante reasons in favour of shortening are stronger than those against.

2.5 Absolute harm

The degree of miserableness also matters when evaluating the extent to which the miserable person has been harmed by having been brought into existence. The notion of harm here is that of absolute rather than comparative harm. This means that the baseline with respect to which harm is determined is specified in terms of an absolute parameter, namely that of a neutral life, and not relative to other possibilities, such as the counterfactual scenario that would have obtained otherwise. The baseline is a normative baseline since it is specified in terms of what the agent has reason to do, in particular what actions the agent has reason to abstain from performing, i.e. what kind of life (namely a miserable life) the agent should not bring into existence. The distance from the normative baseline, and hence the degree of harm, is determined by the strength of the reason the agent had for not bringing about the relevant state of affairs. In other words, the person whose life is miserable has been harmed by having been brought into existence, not insofar as that person has been made worse off relative to non-existence (or relative to some other alternative), but insofar as that person has been put into a state that the agent had reason to not bring into existence (where the strength of the reason corresponds to the degree of miserableness).[30]

This account makes room for non-comparative harms, without there being any non-comparative benefits. The normative baseline is one-sided due to the asymmetry. One is harmed by being brought into existence if one ends up below

[29] We can make sense of unit-comparability in the absence of level-comparability. Since different-number cases involve non-comparability rather than incommensurability, the very same type of value applies.

[30] This notion of absolute harm can be used to make sense of compensation, such that compensating someone for having been brought into existence with a miserable life does not amount to bringing about the status quo ex ante where the person did not exist, but instead amounts to undoing the absolute harm that was imposed on them.

the baseline, whereas one is not benefitted by being brought into existence if one ends up above the baseline. The baseline is one-sided because there is a level, namely that of a neutral life, that is such that no one should be below that level, i.e. bringing someone into existence who is below that level amounts to acting in a way that is contrary to what one has reason to do. This ensures that anyone whose life is miserable (and hence below the level of a neutral life) is harmed by being brought into existence.[31] Yet, there is no level such that everyone (speaking with a possibilist quantifier) should be at or above that level, which implies that no one can be benefitted by being brought into existence. Bringing someone into existence who has a happy life does not constitute a case of doing something that one has reason to do, but merely classifies as a case of not doing something one has reason not to do.[32,33]

3 Conclusion

In the context of a person-affecting approach, the non-comparability of existence with non-existence implies that the fact that a life would be worth living does not give us reasons to bring it into existence. Instead, there are only reasons to improve a life on condition that its existence is given. Since paradigmatically miserable lives are such that they are better the shorter they are, there will be an ordering of shortenings of these lives that is such that, in the limit, they are not lived at all. Bringing these lives into existence, accordingly, violates the shortening consistency constraint. As a result, we have reason not to do so. By contrast, no violation of this constraint is involved in adding happy lives, which means

[31] In probabilistic cases in which it is uncertain whether the life will be lived, i.e. where the life is only to be found in some but not all states of nature, the evaluation of the extent of harm has to be restricted to those states of nature in which the life in question exists, given that the notion of harming is a strong person-affecting notion. This implies that the reason that the agent has to not bring x into existence is a function of the states of nature in which x exists and corresponds to the expected value of x's life.

[32] It might be objected that there cannot be absolute harms without absolute benefits, since not harming someone amounts to benefitting that person. However, failing to subject someone to this type of absolute harm amounts to either failing to create a miserable life by not bringing that person into existence, in which case no one is benefitted since there is no one who can be benefitted (i.e. the subject is missing), or to creating a happy life in which case no one is benefitted in absolute terms since the baseline is one-sided.

[33] If the life that is brought into existence features in several (but not all) alternatives, then the notion of comparative benefit/harm will also be applicable. In such cases, it makes sense to say that x was comparatively benefitted/harmed by being brought into existence in one way rather than another. Here, it is not the absolute level of well-being (whether the life is a happy life or a miserable life) that matters, but the comparative facts. If an agent can either add a life such that the life is slightly miserable or such that it is highly miserable, then bringing about the former possibility classifies as a comparative benefit relative to the latter possibility, but as neither a comparative benefit nor comparative harm relative to the possibility of not bringing the life into existence at all.

that there is no reason not to do so. The asymmetry is thus explained on the basis that the fact that an action violates a consistency constraint renders that action inadmissible and gives us reason not to perform it, whereas the fact that an action satisfies a consistency constraint does not give us any reason to perform that action, but merely makes it the case that it is an admissible action.[34]

References

[1] Arntzenius, F., and McCarthy, D. The two envelope paradox and infinite expectations. *Analysis 57*, 1 (1997), 42–50.

[2] Barrett, J., and Arntzenius, F. An infinite decision puzzle. *Theory and Decision 46* (1999), 101–103.

[3] Heyd, D. *Genethics – Moral Issues in the Creation of People*. University of California Press, 1992.

[4] McMahan, J. Asymmetries in the morality of causing people to exist. In *Harming Future Persons*, M. Roberts and D. Wasserman, Eds. Springer, 2009, pp. 49–68.

[5] McMahan, J. Causing people to exist and saving people's lives. *Journal of Ethics 17* (2013), 5–35.

[6] Narveson, J. Moral problems of population. *The Monist 57*, 1 (1973), 62–86.

[7] Parfit, D. *Reasons and Persons*. Oxford University Press, 1984.

[8] Scott, M., and Scott, A. Infinite exchange problems. *Theory and Decision 57* (2005), 397–406.

[9] Sikora, R. I. Is it wrong to prevent the existence of future generations? In *Obligations to Future Generations*, R. I. Sikora and B. Barry, Eds. Temple University Press, 1978, pp. 112–166.

[34] For helpful comments I would like to thank Dan Halliday, Johan Gustafsson, Kacper Kowalczyk, Theron Pummer, Johann Frick, and especially Tim Campbell as well as audiences at the Conference in Honour of Derek Parfit at the University of Oxford, the Merton Workshop on Applied Formal Ethics, the New Scholarship in Population Ethics conference at Duke University, as well as at Stockholm University and the University of York.