Anonymity and Non-Identity Cases

TOMI FRANCIS

According to the principle of *Anonymity*, any two populations with the same anonymous distribution of wellbeing are equally good. Although it might appear innocuous, Anonymity is a substantive and powerful principle. It implies that differences in the identities of those who exist are evaluatively irrelevant and thereby simplifies population axiology significantly. Perhaps partly for this reason, Anonymity is often assumed in formal treatments of population axiology, sometimes without much in the way of supporting argument.[1] But Anonymity is not obviously true, because it is not obviously true that differences in identity make no all-things-considered evaluative difference.[2] If we are going to accept Anonymity, we had better have a good argument for doing so; my aim in this paper is to provide one.

The crux of my argument for Anonymity is a compelling principle about value in the non-identity cases first discussed by Parfit (1984: ch.17), in which we can either bring about the existence of a better-off person, or alternatively bring about the existence of a different, worse-off person. According to the 'Non-Identity Principle', in such cases, the outcome in which the better-off person comes to exist is better overall. The main technical result of this paper is that given some standard structural assumptions, the Non-Identity Principle implies Anonymity. After demonstrating this result, I shall argue that if the Non-Identity Principle were false, it would follow that replacing worse-off for better-off people, while at the same time improving the lives of existing people, often fails to make things better overall. Since such combinations of changes clearly do make things better overall, the Non-Identity Principle must be true.

Before we turn to the argument from the Non-Identity Principle, let me clarify the background framework in which we shall operate. I take the notion of a life to be primitive, where lives are understood to be individuated by the people living them, and the sorts of lives they are. I assume that there is an 'at-least-as-good-as' relation on lives, which is transitive and reflexive; the 'betterness', 'equal goodness' and other evaluative relations on lives are defined from the at-least-as-good-as relation in the usual way. As shown by Arrhenius (2011), we can identify wellbeing levels with equivalence classes of lives under the equal goodness relation. In practice, I shall represent wellbeing levels by numbers, but this is for ease of presentation only.

1 See for instance Blackorby and Donaldson 1984: 14, Blackorby et al. 2003: 346–47, Broome 2004: 135–36 and Asheim and Zuber 2014: 632.

2 For example, Anonymity conflicts with evaluative versions of McMahan's (2013) 'Weak Asymmetry' between comparative and non-comparative benefits, with Otsuka's (2018) view that it is in itself regrettable for people to be worse off than they might have been, and with various person-affecting views.

Analysis Vol. 81 | Number 4 | October 2021 | pp. 632–639 doi: 10.1093/analys/anab031

The objects of comparison are populations, which I take to be finite sets of lives in which no person appears twice. We are interested in the at-least-as-good-as relation $\succcurlyeq$ on populations, with the derived relations of betterness, equal goodness and so on again being defined in the usual way. I shall allow that $\succcurlyeq$ may be incomplete: that is, there may be populations X and Y, neither of which is at least as good as the other. I shall, however, make two important assumptions about $\succcurlyeq$: I shall assume that it is option-set-independent (in the sense that how two populations compare does not depend on any underlying option set in which they are considered) and also transitive.[3] Although I believe that we should accept choice-set-independence and transitivity, I shall not argue for them in this paper. Those who doubt one or both of these assumptions, and also find Anonymity to be incredible, might take the lesson of this paper to be that we had better reject transitivity or option-set-independent betterness in order to maintain the Non-Identity Principle without incurring a commitment to Anonymity.

Anonymity, stated more precisely, is the claim that if two populations are *Anonymously Equivalent*, meaning that we can find a bijection between the lives in each of the two populations which preserves wellbeing levels, these two populations are equally good.

That is enough in the way of background; on to the argument from the Non-Identity Principle to Anoynmity. Consider a standard non-identity case, in which one must choose between bringing Eve into existence, with a better life, or instead bringing Adam into existence, with a worse (but still good) life. Imagine that nobody else is affected by this choice, and that there are no other potentially morally relevant differences at play: Adam and Eve are equally deserving, no impersonally valuable things are affected, and so on. (A *ceteris paribus* clause like this should be taken to apply to all further cases and principles appealed to in this paper.)

One thing I think we should say about this case is that it would be better overall for Eve to exist with the better life, rather than for Adam to exist with the worse life. Note that this claim is purely evaluative. I am not claiming that it would be wrong to bring Adam into existence, or even that we would have moral reason to bring Eve into existence rather than Adam. These last two claims, while plausible, are somewhat more controversial than the evaluative claim.[4]

It also seems that it would be better overall for Eve to exist rather than Adam regardless of the number, identities and wellbeing levels of other people,

3 The option-set-independence of betterness is disputed by Roberts (2003, 2011), Voorhoeve (2013) and Frick (2014, 110–20). The rejection of transitivity is discussed at length by Temkin (2012) and is argued for by Temkin (1987, 1996) and Rachels (1998, 2001, 2004). One standard defence of transitivity can be found in Broome 2004: 50–51.

4 Boonin (2014: 151) provides an influential defence of the view that we have no moral obligation to bring about better lives in non-identity cases. But even he admits that in his model non-identity case, 'conceiving [the better off person] would produce significantly better consequences'.

provided that these people are entirely unaffected by the choice. More generally, it seems that we should accept the

Weak Non-Identity Principle

> If life l_1 is better than life l_2, then for any population X not including the l_1 or l_2 persons, $X + l_1$ is better than $X + l_2$.

The Weak Non-Identity Principle supports the stronger claim that, additionally, if lives l_1 and l_2 are equally good, then $X + l_1$ is equally as good as $X + l_2$. This is because the Weak Non-Identity Principle implies that the balance between the populations $X + l_1$ and $X + l_2$ is sensitive to arbitrarily small improvements and deteriorations: it would be better to add a slightly better version of l_1 than to add l_2, and vice versa. This kind of sensitivity is widely (and correctly) held to be strong evidence of equal goodness; conversely, insensitivity to small improvements is 'the mark of incommensurability' (Raz 1986: 325–26) or symptomatic of 'imprecise equality' (Parfit 2016: 115). Broome (2004: 21) goes so far as to *define* equal goodness to hold whenever we have sensitivity to all improvements/deteriorations. We are thus justified in inferring from the Weak Non-Identity Principle the

Strong Non-Identity Principle

> Life l_1 is at least as good as life l_2 if and only if for any population X not including the l_1 or l_2 persons, $X + l_1$ is at least as good as $X + l_2$.

(Henceforth just the 'Non-Identity Principle'.)

We shall now see that the Non-Identity Principle implies Anonymity. First note that the Non-Identity Principle requires that if a single person in one population is replaced by a different person at the same wellbeing level, the resulting population will be equally good. We can use this fact to show that Anonymously Equivalent populations must be equally good when they are disjoint (by which I mean that no person belongs to both populations). Given two disjoint, Anonymously Equivalent populations, one population can be transformed into the other by simply replacing each person by their same-wellbeing counterpart along a bijection witnessing Anonymous Equivalence, one at a time. The Non-Identity Principle implies that each replacement results in an equally good population, and transitivity then yields that the first population is equally as good as the last.

This suffices to show that any two disjoint, Anonymously Equivalent populations must be equally good. Now suppose we have two arbitrary Anonymously Equivalent populations, A and B, which might not be disjoint. We may construct a third population C, which is Anonymously Equivalent to, and disjoint with, both A and B. Since each of A and B are equally as good as C by the previous argument, transitivity requires that A and B are equally as good as each other. The Non-Identity Principle therefore implies Anonymity, given transitivity. In fact, something stronger can be proved: the Non-Identity Principle implies the principle of *Anonymous Pareto*, according to which a population is better than

another when the first is Anonymously Equivalent to a population which is weakly pareto superior to the second (i.e. better for some, and at least as good for all). The proof of this claim uses the same sort of idea as the proof of Anonymity: we replace each person by their better-off or equally well-off counterpart one at a time, and use transitivity to chain the betterness claims together.

On the face of it, it is surprising that the Non-Identity Principle implies Anonymity. The Non-Identity Principle only says that what we may call 'non-identity improvements' – replacements of worse-off people by different better-off people – make the world better. But the Non-Identity Principle does *not*, on the face of it, say that non-identity improvements make just as much evaluative difference as the provision of ordinary comparative benefits. The Non-Identity Principle thus looks consistent with a 'Two-Tier View' (Parfit 2017), on which non-identity improvements matter, but they matter less, evaluatively speaking, than ordinary comparative benefits. Yet the Two-Tier View implies that the Non-Identity Principle is true and Anonymity is false, so how could the Non-Identity Principle imply Anonymity? The answer is that, as Parfit shows, the Two-Tier View is cyclic (or option-set-dependent), and therefore intransitive (or option-set-dependent). One might nevertheless wonder whether there might be some transitive, option-set-independent population axiology which satisfies the Non-Identity Principle but not Anonymity. The preceding argument confirms that there is no such population axiology.

I believe most readers will find the Non-Identity Principle rather plausible. But perhaps some will find it more palatable to reject the Non-Identity Principle than to accept Anonymity. Can anything be said against this position? Yes. The consequences of rejecting the Non-Identity Principle are more implausible than it might at first appear. At least, this is true if we accept plausible principles for comparing same-person populations, such as the principle of *Same-Person Anonymity*, according to which Anonymously Equivalent populations are equally good if they contain exactly the same people. Same-Person Anonymity is the uncontroversial part of Anonymity. It implies no contentious claims about how the importance of benefiting people compares to the importance of replacing people; it merely states that it does not make an evaluative difference who gets what, when everyone exists regardless of which population comes about.[5] It is a very plausible principle, so I

5 It might be claimed the Same-Person Anonymity is unacceptably aggregative. Consider the following case, due to Parfit (2003: 383, n.16), in which one faces a choice between two populations involving the same one hundred people:

(i) Person One is at level 1, Person Two at level 2, and so on.

(ii) Person One is at level 100, and everyone else is one level lower down.

Same-Person Anonymity implies that (i) and (ii) are equally good, whereas one might think that (ii) is better, since (i) involves a large loss for Person One in return for minor gains for the others. However, we can replace Same-Person Anonymity with the non-aggregative principle of Pairwise Same-Person Anonymity, which is the restriction of Same-Person Anonymity to cases in which the identities of at most two people are permuted. Given our background assumptions, Pairwise Same-Person Anonymity is equivalent to Same-Person Anonymity.

shall take it for granted in what follows (and in any case, the following argument can be adapted so as to make do with different principles for comparing same-person populations).

Suppose that the (Weak) Non-Identity Principle is false, so that for some population A, a better life l_1 and a worse life l_2, $A + l_1$ is not better than $A + l_2$. We shall assume that A contains a life at some wellbeing level w strictly between the l_1 and l_2 levels; unless the Non-Identity Principle fails only when this is not the case (which would be bizarre), this assumption is innocent. Write A' for a population that differs from A only in that one person at w is now at the wellbeing level of l_1, and write l'_1 for a life in which the l_1 person has wellbeing level w. By assumption, $A + l_1$ is not better than $A + l_2$. Same-Person Anonymity implies that $A + l_1$ is equally as good as $A' + l'_1$. Given transitivity, these two claims imply that $A' + l'_1$ is not better than $A + l_2$. But this would mean that it is sometimes not better overall for existing people to be made better off, while at the same time worse-off people are replaced by different better-off people. That is an implausible conclusion.

It is easier to see what is going on, and why the conclusion is implausible, if we consider a more concrete case. Suppose that, because we reject the Non-Identity Principle, we believe that A is not better than B in the case illustrated by the table below, where Ω represents non-existence:

	Adam	Steve	Eve
A	60	80	Ω
B	60	Ω	40
C	80	60	Ω

Intuitively, C is better than B. Same-Person Anonymity implies that A and C are equally good. Transitivity then implies that A is better than B, in line with the Non-Identity Principle. If we want to reject this instance of the Non-Identity Principle, it thus looks like we have to say that C is not better than B, despite appearances. But compare B and C directly. We can imagine going from C to B by making two changes. First, we make Adam worse off by twenty units of wellbeing. Second, we replace Steve with the worse-off Eve. The first change clearly makes the world worse overall.[6] In this particular case, it is

6 This claim might be false on a suitably extreme egalitarian view, on which decreases in wellbeing are not worse overall when they sufficiently reduce inequality. I think that views like this are not particularly plausible, but it is worth noting that if we accept such a view, it might follow that the first change does not make things worse overall, but on the other hand it would seem that the second change does make things worse overall, because it brings us from an equal to an unequal population. So on this sort of view we would face a similarly compelling argument to the effect that the combined change must be for the worse, since if we combine a bad change with a change that only makes someone worse off, the combined change presumably makes things worse overall.

implausible to claim that if the second change occurs as well, the world is not made worse overall.[7] To avoid this implausible claim, we should accept the Non-Identity Principle.

The preceding argument is even more forceful in cases in which multiple worse lives could be replaced by multiple better lives. If the Non-Identity Principle were false, so that replacements of worse by better lives sometimes fail to make the world better overall, it would seem natural to think that replacements of many worse by many better lives would likewise fail to make the world better overall. But this view has absurd implications. Consider two possible futures. In the first future, the present seven billion inhabitants of Earth will each enjoy 80 units of wellbeing, and there will exist seven billion future individuals, each with 60 units of wellbeing. In the second future, the present inhabitants will instead have 60 units of wellbeing, while seven billion future individuals, who are non-identical to the individuals existing in the first future, will have only 40 units of wellbeing. If we suppose that it does not make the world better to replace any number of lives at level 40 by the same number of lives at level 80, an argument exactly analogous to that of the preceding paragraph shows that, given Same-Person Anonymity and Transitivity, the second future is not worse than the first. But we should not accept this conclusion. A change from the first to the second future makes every present person much worse off, makes seven billion future people much worse off than the seven billion who would otherwise have existed, and affects nobody else. This is obviously a change for the worse, and we should reject any principle which says otherwise.

Let us take stock. Suppose we accept the transitivity and choice-set-independence of betterness. We must then accept the Non-Identity Principle, on pain of absurd conclusions. Since the Non-Identity Principle implies Anonymity, we must accept Anonymity as well. Doing so amounts to substantial progress. There are, broadly speaking, two basic questions in population axiology. The first is the question of how to compare populations which differ with respect to the number of people they contain. The second is the question of how to compare populations which differ with respect to the identities of the people they contain. Anonymity provides a full answer to

7 There are similarities between my argument here and Broome's (2004: 169–70; 2005: 202–8) 'greediness' objection to the Intuition of Neutrality. Some readers might worry that criticisms of Broome's argument, such as those levelled by Rabinowicz (2009) or Frick (2017), might apply equally to my own argument. But this is not the case. My argument does not appeal to the general claim that a bad thing plus a neutral thing must be bad, or that a bad thing plus a thing which is not bad must be bad. It instead appeals to the more specific claim that replacing better-off people with different worse-off people cannot swallow up the badness of making existing people worse off. This specific claim is plausible even if the more general claims are false.

the second question, namely that such differences in the identities of people make no all-things-considered evaluative difference.[8]

Funding

This work was supported by an Arts and Humanities Research Council Doctoral Training Partnership [grant number AH/L503885/1] and by a Derek Parfit Memorial Prize from the Global Priorities Institute, University of Oxford.

St John's College, University of Oxford, UK
tomi.francis@philosophy.ox.ac.uk

References

Arrhenius, G. 2011. The impossibility of a satisfactory population ethics. In *Descriptive and Normative Approaches to Human Behaviour*, eds. E. N. Dzhafarov and L. Perry, 1–26. Singapore: World Scientific Publishing Co.

Asheim, G.B. and S. Zuber. 2014. Escaping the repugnant conclusion: rank-discounted utilitarianism with variable population. *Theoretical Economics* 9: 629–50.

Blackorby, C., W. Bossert and D. Donaldson. 2003. The axiomatic approach to population ethics. *Politics, Philosophy & Economics* 2: 342–81.

Blackorby, C. and D. Donaldson. 1984. Social criteria for evaluating population change. *Journal of Public Economics* 25: 13–33.

Boonin, D. 2014. *The Non-Identity Problem and the Ethics of Future People*. New York: Oxford University Press.

Broome, J. 2004. *Weighing Lives*. Oxford: Oxford University Press.

Broome, J. 2005. Should we value population? *Journal of Political Philosophy* 13: 399–413.

Frick, J. 2014. *'Making People Happy, Not Making Happy People': A Defense of the Asymmetry Intuition in Population Ethics*. PhD thesis, Harvard University.

Frick, J. 2017. On the survival of humanity. *Canadian Journal of Philosophy* 47: 344–67.

Frick, J. forthcoming. Context-dependent betterness and the mere addition paradox. In *Ethics and Existence: The Legacy of Derek Parfit*, eds. J. McMahan, T. Campbell, J. Goodrich and K. Ramakrishnan. Oxford: Oxford University Press.

McMahan, J. 2013. Causing people to exist and saving people's lives. *The Journal of Ethics* 17: 5–35.

Otsuka, M. 2018. How it makes a moral difference that one is worse off than one could have been. *Politics, Philosophy & Economics* 17: 192–215.

8 For helpful comments and discussion on earlier drafts of this paper, I would like to thank Hilary Greaves, Petra Kosonen, Kacper Kowalczyk, Michal Masny, Teruji Thomas, Elliott Thornley, two anonymous referees and the participants of the February 2020 Stockholm Workshop on Population Ethics and the Non-Identity Problem. I owe special thanks to Todd Karhu and Johan E. Gustafsson for their insightful comments on multiple drafts of this paper.

Parfit, D. 1984. *Reasons and Persons*. Oxford: Clarendon Press.

Parfit, D. 2003. Justifiability to each person. *Ratio* 16: 368–90.

Parfit, D. 2016. Can we avoid the repugnant conclusion? *Theoria* 82: 110–27.

Parfit, D. 2017. Future people, the non-identity problem, and person-affecting principles. *Philosophy & Public Affairs* 45: 118–57.

Rabinowicz, W. 2009. Broome and the intuition of neutrality. *Philosophical Issues* 19: 389–411.

Rachels, S. 1998. Counterexamples to the transitivity of *better than*. *Australasian Journal of Philosophy* 76: 71–83.

Rachels, S. 2001. A set of solutions to Parfit's problems. *Noûs* 35: 214–38.

Rachels, S. 2004. Repugnance or intransitivity: a repugnant but forced choice. In *The Repugnant Conclusion: Essays on Population Ethics*, eds. J. Ryberg and T. Tännsjö, 163–86. London: Kluwer Academic.

Raz, J. 1986. *The Morality of Freedom*. Oxford: Clarendon Press.

Roberts, M.A. 2003. Is the person-affecting intuition paradoxical? *Theory and Decision* 55: 1–44.

Roberts, M.A. 2011. The asymmetry: a solution. *Theoria* 77: 333–67.

Temkin, L.S. 1987. Intransitivity and the mere addition paradox. *Philosophy & Public Affairs* 16: 138–87.

Temkin, L.S. 1996. A continuum argument for intransitivity. *Philosophy & Public Affairs* 25: 175–210.

Temkin, L.S. 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford: Oxford University Press.

Voorhoeve, A. 2013. Temkin's critique of transitivity. *Economics and Philosophy* 29: 409–23.

A note on the Wilhelmine Inconsistency

JON ERLING LITLAND

1. Introduction

Wilhelm (2021) has recently shown that certain widely accepted grounding principles are inconsistent with standard principles of propositional identity; call this the Wilhelmine Inconsistency. The proof of the inconsistency is unimpeachable, but what is its import? This note argues that – unlike the puzzles of ground (Fine 2010) and the Russell-Myhill paradoxes for ground (Fritz forthcoming) – the Wilhelmine Inconsistency should not worry grounding enthusiasts. This is established by sketching two natural ground-theoretic accounts of propositional identity, and showing that on one account a grounding principle is clearly incorrect, and that on the other account the principles of propositional identity are clearly incorrect. Both accounts are conservative in

Analysis Vol. 81 | Number 4 | October 2021 | pp. 639–647 doi: 10.1093/analys/anab048