Roberts on Depletion: How Much Better Can We Do for Future People?

MARK E. GREENE
University of Delaware

Suppose that Depletion will reduce the well-being of future people. Many of us would like to say that Depletion is wrong because of the harm to future people. However, it can easily be made to seem that Depletion is actually harmless – this is the non-identity problem. I discuss a particularly ingenious attempt by Melinda Roberts to attribute a harm to Depletion. I will argue that the magnitude of Roberts's harm is off target by many orders of magnitude: it is just too tiny to explain the intuitive wrong of Depletion.

THE NON-IDENTITY PROBLEM

Consider:

> *The Simple Policy Choice*: We must choose between Conservation and Depletion. If we choose Conservation we will enjoy high levels of well-being now and so will those who live seven generations from now. If we choose Depletion, we will be even better off, enjoying *very* high levels of well-being, but members of the seventh generation will be left with only moderate levels of well-being.[1] The options and results are summarized in Table 1 below.

I hereby stipulate that Conservation and Depletion are the only two options, that only the present and seventh generations will be affected, that there is no variation in well-being among members of the same generation and that none of this makes any difference to the number of people who will exist.

From our point of view there is obvious appeal to Depletion, so why not make merry now? Obviously, one might think, because those alive seven generations hence will enjoy high levels of well-being if we are prudent, but will plod along with merely moderate levels of well-being if we squander resources. Surely, the killjoys will insist, the cost to the seventh generation carries at least some ethical weight against considerations of present self-interest.

There is a problem for the killjoys. Major policy decisions shift economic activity and those shifts will change who meets, mates and

[1] This is a variation of Parfit's case; Derek Parfit, *Reasons and Persons* (Oxford, 1984), pp. 361–62.

© Cambridge University Press 2015
Utilitas Vol. 28, No. 1, March 2016
doi:10.1017/S095382081500031X

CrossMark
https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

Table 1. *The simple policy choice: options and results*

	Conservation	Depletion
Present Generation	High well-being	Very high well-being
Seventh Generation after Us	High well-being	Moderate well-being

procreates. Because different present couplings lead to different future offspring, the people who will be alive seven generations from now will owe their existence to our present policy choice. If the post-Depletion people will owe their existence to Depletion, it is hard to see how that policy will harm them. If the post-Depletion people will not be harmed by Depletion, it is hard to think of anyone else who will be. Finally, if no one will be harmed by Depletion, it is hard to see what could be wrong with it. To those of us who can't shake off the feeling that there really is something seriously problematic about Depletion, this is the non-identity problem.[2]

I will discuss a particularly ingenious attempt by Melinda Roberts to attribute a harm to Depletion.[3] I think her project is the right one: namely, to identify a harm done by Depletion to denizens of the seventh generation. However, to answer the substance of the non-identity problem, the harm of Depletion must carry significant ethical weight when balanced against the benefits of Depletion for the present generation. I will argue that, even if Roberts has detected a harm in Depletion, the magnitude of that harm is off target by many orders of magnitude: the harm is just too tiny to give the moral scales more than a negligible nudge in favour of Conservation.

COULD WE HAVE DONE BETTER FOR FUTURE PEOPLE?

Melinda Roberts points out a feature of cases like Depletion which, while never in dispute, is usually overlooked. The claim that big policy changes will change who meets, mates and procreates is a probabilistic one: it is unfathomably unlikely that the same couplings and conceptions would occur in the sequelae to both Conservation and Depletion, but it is neither biologically nor metaphysically impossible.

This might not seem to make much difference. Consider the thinking of a denizen of the seventh generation post-Depletion: call him George Jetson. Suppose that George's actual well-being is 5 hedons. Feeling dissatisfied with life, George calculates that he would have enjoyed 50 hedons of well-being if he had lived post-Conservation.

[2] Parfit, *Reasons and Persons*, pp. 351–80; Michael D. Bayles, 'Harm to the Unconceived', *Philosophy & Public Affairs* 5 (1976), pp. 292–304.

[3] Melinda A. Roberts, 'The Non-Identity Fallacy: Harm, Probability and Another Look and Parfit's Depletion Example', *Utilitas* 19 (2007), pp. 267–311.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

But then he realizes he almost certainly would not have existed post-Conservation. Even if George's post-Conservation chances of coming into existence were as implausibly high as one in a trillion, 50 hedons times 1/1,000,000,000,000 is only 50 picohedons of expected post-Conservation well-being. George grudgingly concludes that he's better off with what he's got.

George is attempting to identify a personal harm due to Depletion by comparing his post-Depletion well-being with that he would have enjoyed post-Conservation. George evaluates his post-Conservation well-being on the basis of an expected value calculation, which includes the prior uncertainty of his coming into existence. But George does not consistently apply an expected value approach when it comes to evaluating his post-Depletion well-being. Instead, he makes this evaluation assuming his knowledge of how things actually turned out, taking his existence as a given. Roberts argues that George errs by 'drawing haphazardly from a potpourri of *actual* and *expected* well-being levels'[4] and that his 'comparisons between actual and expected well-being must be recognized as unreliable indicators of harm.'[5]

According to Roberts, it is possible to identify a harm done by Depletion using either an actual value (objective) approach[6] or an expected value approach.[7] The key is to apply one approach or the other consistently and not to mix the two. On an actual value approach, Roberts sees a harm to George in the shortfall of his actual, moderate level of post-Depletion well-being as compared to a post-Conservation possibility in which George exists and enjoys high well-being. Roberts, however, does acknowledge implications of her actual value approach that, as she delicately puts it, 'may seem implausible'.[8] She also acknowledges that her actual value approach must tie harm to an outcome in which George exists 'despite the fact that that outcome is one the agent in a million years could never achieve'.[9] This disconnect between what counts as a harm and what an agent can reasonably aim to achieve undermines the action-guiding function that one would wish one's ethical considerations to provide.

Doubts about the plausibility and applicability of an actual value approach motivate turning to an account that 'would make harm the

[4] Melinda A. Roberts, 'The Nonidentity Problem and the Two Envelope Problem: When Is One Act Better for a Person Than Another?', *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, ed. Melinda A Roberts and David Wasserman (Dordrecht, 2009), pp. 201–28, at 204.

[5] Roberts, 'The Non-Identity Fallacy', p. 310.

[6] Roberts, 'The Non-Identity Fallacy', p. 277.

[7] Roberts, 'The Non-Identity Fallacy', p. 282.

[8] Roberts, 'The Non-Identity Fallacy', p. 279.

[9] Roberts, 'The Non-Identity Fallacy', p. 280.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

kind of thing we often know how – step by step – to avoid imposing, should we chose to avoid imposing it.'[10] When we focus on what agents know at the time of their decision to act, our attention turns away from George's retrospective view and back to the perspective of present policymakers. As they face the choice between Depletion and Conservation, policymakers do not know how things will actually turn out but they can make assessments of expected value. Roberts notes that any consistent evaluation of expected well-being 'must be restricted to information the agent is in a position to grasp at that critical time prior to the performance of the act under scrutiny'.[11] This is right. Suppose I steal £100,000 of your retirement savings and buy enough lottery tickets for a one in a million chance of the £20 million mega-jackpot. Against the odds, I do win and I give you half the pot. By sheer dumb luck it all worked out and you will have no regrets. But your retrospective lack of regret does not justify stealing. Assuming you'd get £10 million in the event I won and nothing otherwise, I reduced your expected wealth from £100,000 to £10 at the time I, eh-hem, reinvested your savings. That was wrong of me.

Whether we pursue Depletion or Conservation, George Jetson's present expected well-being depends on just two things: first, on his present odds of existing given our policy decision; second, on the well-being he will enjoy if he does exist in the aftermath of that decision. Since George's present expected well-being post-Depletion depends, in part, on his odds of existing, we cannot just suppose that Depletion will lead to George's existence. Indeed, from the perspective of the present generation, the existence of any particular possible future person is highly unlikely on any policy. Now, it may be that present policy choices do impact the objective probability of George coming into existence. However, we have no prior knowledge of what that impact is and, thus, nothing to motivate any difference between our present subjective probability of George existing post-Depletion and that of his existing post-Conservation. Therefore, if our policy choice has any impact on George's present expected well-being, it cannot be because of the difference it makes to his odds of existing, it can only be because of the difference it makes to the level of well-being he will enjoy if he does exist. That difference of well-being, of course, counts in favour of Conservation.[12]

[10] Roberts, 'The Non-Identity Fallacy', p. 279.

[11] Roberts, 'The Non-Identity Fallacy', p. 297.

[12] Roberts argues this rather differently, via a lot of work aiming to show that present agents can do things to make the odds of some future person existing under Conservation as high as they are under Depletion. I think she is not successful in this. I also think she does not need to be because the epistemic limitations of the agent are enough to equalize the subjective probabilities of any particular person existing following either policy.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

VANISHINGLY SMALL HARMS

If Roberts has been successful, she has shown how a consistently applied expected value approach can support the claim that Depletion does harm future people, despite the non-identity problem: Depletion reduces the present expected well-being of particular possible future people. I have two layers of doubts about this strategy that I will not pursue here. First, I doubt that sense can be made of the present odds of any particular future person existing. Second, if we do manage to make sense of the odds, I doubt that they will be representable by real numbers. Although there will only be finitely many *actual* people in any future generation, I suspect that there are infinitely many *possible* future people and that the odds of any particular one of them existing are infinitesimal. The challenge of calculating expected well-being using surreals is rather daunting.

Setting aside these two concerns by assuming there are finitely many possible future people and that any possible future individual's odds of existence are real, I remain suspicious that the harm Roberts sees in Depletion will be too tiny to make much of a dent in the substance of the non-identity problem. Although the letter of the non-identity challenge may be answered by showing that there is some harm rather than no harm in choosing Depletion, the substance of the problem is to justify the intuition that Depletion's impact on the well-being of future people gives present people a substantial moral reason at least to consider forgoing the benefits of Depletion. Bearing in mind Roberts's admonition that the expected value approach to our present policy decision 'must be restricted to information the agent is in a position to grasp at that critical time prior to the performance of the act under scrutiny',[13] present policymakers must weigh the impact of Depletion on the expected well-being of the present generation against its impact on the expected well-being of the seventh generation. If we do choose Depletion rather than Conservation, members of the present generation will enjoy quite a large well-being advantage. Since this well-being advantage is certain to occur (by stipulation), Depletion confers quite a large advantage of expected well-being on members of the present generation. Even if Roberts has managed to identify some harm in terms of Depletion's impact on the expected well-being of members of the seventh generation, my suspicion is that the expected harm to the seventh generation is sure to be insignificant when weighed against the large expected benefit to the present generation.

To substantiate my suspicion I will show that Conservation will very clearly yield only a negligible total expected welfare advantage

[13] Roberts, 'The Non-Identity Fallacy', p. 297.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

for the seventh generation. The total is negligible even I exaggerate the advantages of Conservation over Depletion by inflating the welfare impact of the policy choice, by aggregating over ridiculously high estimates of the total population of the seventh generation and by grotesquely overestimating the chances of a particular future person coming into existence.

Let's start with the well-being impact of the policy choice. I stipulated that people in the seventh generation would have good lives post-Conservation but only OK-ish lives post-Depletion. I will now inflate that difference and stipulate that, post-Conservation, every member of the seventh generation will live for a hundred years at a steadily great level of well-being – think of the happiest person you know bubbling along at their most slappably chirpy for a century straight. On the Depletion side, let's stipulate that all of that is gone and every moment of the life of a post-Depletion seventh generationer might as well have been spent anaesthetized for all the well-being it affords. For convenience, let's call each life year of great well-being a GWLY: then each seventh generationer will get 100 GWLYs if we choose Conservation but 0 GWLYs if we choose Depletion.

Now for the population of the seventh generation. Just to be silly, let's suppose that the population increases tenfold with each generation. To keep the numbers nice and round, let's set the present generation at ten billion. The first generation after ours will then have 100 billion people, the second a trillion people, and so on until the seventh generation contains one hundred quadrillion people; that's 100 million billion, or 10^{17}, or about 200 people per square metre of the Earth's surface – land and sea. I don't think I can be fairly accused of underestimation.[14]

I really don't know how to go about estimating the present odds of a particular individual existing in the seventh generation. Luckily, all I need is a number that's so plainly a crazy overestimate that no one who gives a shred for their integrity will accuse me of lowballing. So let's pretend that the existence of a particular target person requires only the fertilization of an egg (any egg will do) with the right particular sperm from the right particular father. Let's also pretend that, if the right father exists then the right sperm is guaranteed to be present in an insemination. Finally, let's pretend that an insemination containing the right sperm is guaranteed to produce the target person as long as it is the right sperm that gets there first and fertilizes the egg. Now, typically, there are about 200 million sperm per ejaculation of which over half show progressive mobility and over 10 per cent have normal

[14] Thanks to Wolfram Alpha (www.wolframalpha.com) for these estimates and doing most of the actual calculations that follow.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

morphology,[15] but, to be on the safe side, I'll assume that the odds of the right sperm being the one that fertilizes the egg are as high as one in a million.

Putting all this together, for George Jetson to exist in the seventh generation requires, first, that the right five greats grandfather exists in the present generation and sires the right four greats grandfather. Assuming the right five greats grandfather exists, we can rely on all the silly assurances about the right insemination taking place to put the odds of the right four greats grandfather being born at an indefensibly high one in a million. To get all the way to George in the seventh generation, the process must be successfully repeated seven times. That's seven rolls of a million-sided reproductive die. The odds of succeeding at all seven steps are one millionth to the power of seven, or 10^{-42}.

Now we can work out a number for George's present expected well-being if we choose Conservation. The pay-off is 100 GWLYs, which is 3,155,760,000 GWLSs (Great Well-being Life Seconds). But the probability of George existing to enjoy the benefit is 10^{-42}, so the present expected benefit to George of Conservation is $10^{-42} \times 3{,}155{,}760{,}000 = 3.15576 \text{ x } 10^{-33}$ GWLSs. Even if Depletion would rob George of all of that, three decillionths of a second of happiness does not amount to a whole hill of welfare beans.

Not to worry, there are many more people than George in the seventh generation. This suggests a route to recognizing a larger gap between the expected well-being of a post-Conservation seventh generation and that of a post-Depletion seventh generation. Instead of looking only at the impact on George's expected well-being, someone in search of a bigger hit to expected well-being might suggest aggregating expected well-being for every member of the seventh generation. Having stipulated that there is no variation in well-being within a generation, we can estimate aggregate expected well-being simply by multiplying the estimate for George by the total seventh generation population of 100 quadrillion. How much bigger does this make the well-being gap? Well, 100 quadrillion times $3.15576 \text{ x } 10^{-33}$ gives eminently negligible total of $3.15576 \text{ x } 10^{-16}$ GWLSs. What the compounding population gave, the compounding improbability of existence hath taken away. Even if Roberts has drawn thirty quadrillionths of the bite of the non-identity problem, that's just not enough to count as a solution.

[15] Trevor G Cooper, Elizabeth Noonan, Sigrid von Eckardstein, Jacques Auger, H. W. Gordon Baker, Hermann M Behre, Trine B Haugen, Thinus Kruger, Christina Wang, Michael T Mbizvo and Kirsten M Vogelsong, 'World Health Organization Reference Values for Human Semen Characteristics', *Human Reproduction Update Human Reproduction Update* 16 (2010), pp. 231–45, at 237.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

ALL POSSIBLE SEVENTH GENERATIONERS

To get a larger total present harm, we could try to aggregate over a larger population. For every actual conception there will be myriad others that could happen but never do and correspondingly many future people who are now possible but will never be actual. Perhaps a more noticeable harm will result from aggregation, not just over *actual* seventh generationers, but over all *possible* seventh generationers. There is some justification for doing this. Although we cannot now know whether George Jetson will turn out to be actual, this does not make any obvious difference to the calculation of his present expected well-being: his minuscule odds of becoming actual are already part of the calculation. And there's nothing special about George. He is but one representative of the many who are, from the present perspective, possible denizens of the seventh generation. So, if we can assess the impact of Depletion on George's present expected well-being without regard to whether he will ever actually be born, the same goes for all other presently possible seventh generationers. We can then aggregate over all of them in the firm expectation of reaching a total impact far higher than that attainable by aggregation over the actual 100 quadrillion.

Unfortunately, the results of aggregation over all possible future people are highly implausible. Consider:

> *The Even Simpler Policy Choice*: We must choose between Conservation and Depletion. The choice will affect only the present generation and the next generation. On either policy, 10 billion people will live in each generation and every one of them will live for 100 years. Everyone will have great well-being every day of their lives except for those days spent in departmental meetings. Departmental meeting days contribute nothing to anyone's lifetime well-being. If we choose Conservation, every member of the present generation and every member of the next generation will have to devote one day in ten to departmental meetings so, over a 100-year lifetime, 10 years will be sunk into departmental meetings leaving 90 life years of great well-being (90 GWLYs) per person. If we choose Depletion, we will never have to sit through a departmental meeting but the next generation will have to spend every other day immersed in *Robert's Rules of Order*. The options and results are summarized in Table 2 below.

If we are to aggregate present expected harm over all possible members of the next generation, then this case and others like it are woefully under-described: we need more information to work out the

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

Table 2. *The Even Simpler Policy Choice: options and results in GWLYs per person*

	Conservation	Depletion
Present Generation	90	100
Next Generation	90	50

number of *possible* next generationers so that we can aggregate over all of them. Suppose that we can make the estimate by first working out the total number of possible children for each present woman and then calculating the sum for all 5 billion present women. A woman who has no eggs can have no children, so she will add nothing to the total of possible next generationers. A woman with one viable egg can have a child, so that one egg in that one woman adds to the number of possible people in the next generation: there's no need to worry how many possible children are added by one viable egg, as long as we continue to assume that the number is finite. What about a woman with two or more viable eggs? Dropping the earlier, silly pretence that any egg will do, it is most implausible that the set of possible children who might develop from one egg will exactly match the set that might develop from a distinct egg and, therefore, additional eggs will add to the woman's total number of possible offspring. With all this in mind, we can set up two variants of the Even Simpler Policy Choice:

> *Even Simpler Policy Choice – Spendthrift Ovaries*: Each present woman has a million viable eggs, two of which are fertilized to produce two members of the next generation.
> *Even Simpler Policy Choice – Efficient Ovaries*: Each present woman has two viable eggs, both of which are fertilized to produce two members of the next generation.

If the appropriate aggregation really is over all possible next generationers, then these biological contingencies are morally relevant because they change the number of possible next generationers. Assuming that each additional egg contributes equally to the total of possible next generationers, the total present expected harm of Depletion to the next generation will be half a million times greater on the Spendthrift Ovaries variant than on the Efficient Ovaries variant of the Even Simpler Policy Choice. Here are a few corollaries of this, none of which are flattering to the idea of aggregating over possible people. First, since the variants are orders of magnitude apart, they

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

cannot both produce aggregates that are plausible estimates of the intuitive harm of Depletion. Second, the strategy implies that the efficiency of human ovaries is morally relevant. Third, if we are going to choose Depletion, we can greatly mitigate the expected harm if we can give present women efficient ovaries and we can enhance the expected benefits of Conservation by minimizing ovarian efficiency. Fourth, to generalize the third corollary, if we have moral reason to mitigate harms and enhance benefits, then we have moral reason to increase the efficiency of present ovaries when future people are expected to be harmed and to reduce their efficiency when future people are expected to benefit. All of these corollaries are highly implausible.

THE PROBLEM REMAINS

At first glance, it seemed obvious that Depletion would significantly harm future people, but at the heart of the non-identity problem is an argument that Depletion is actually harmless. More generally, there is a problem when some plausible calculation of future harm is significantly mismatched with the intuitive magnitude of the harm and with the seriousness of the corresponding wrongdoing. The magnitude of Depletion's expected harm to future people matters because that harm must be weighed against the significant expected cost of Conservation for present people and, in this context, harms of the order of a few tens of quadrillionths of a GWLS cannot explain the intuitive moral problem with Depletion. Some respond by trying to ground the intuitive wrong of Depletion in something other than harm,[16] but this raises a question that I've never seen answered: why does the seriousness of the alleged harmless wrongdoing just happen to approximate the seriousness of the initially intuitive, harm-based wrongdoing? If there is a match it seems like gerrymandering but, if not, the intuitive mismatch underlying the non-identity problem is untouched.

These are some of the reasons for which I have great sympathy with Melinda Roberts's project of accounting for wrongs to future people in terms of harms to future people. However, Roberts's approach to identifying a harm in terms of the present expected well-being of future people fails to resolve the mismatch between the magnitude of the harm identified and that of the intuitive harm and wrongdoing. Even if we aggregate the expected harm over vast populations of actual future people, the harm is still so footling that it might as well be zero for all

[16] Gregory Kavka is a prominent proponent of harmless wrongdoing. Gregory S. Kavka, 'The Paradox of Future Individuals', *Philosophy and Public Affairs* 11 (1982), pp. 93–112.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press

practical purposes. Alternatively, if we aggregate over possible future people, the harm of Depletion varies by orders of magnitude depending on morally irrelevant biological contingencies. None of these results is significantly less jarring than the argument for the harmlessness of Depletion that got into the non-identity problem in the first place.[17]

mkgreene@udel.edu

[17] I am very grateful to Chris Boorse, Richard Hanley, Jeff Jordan, Jonathan Justice, Tom Powers, Joel Pust, Fred Schueler and Seth Shabo for conversation and comments on earlier versions of this article. I would also like to thank conference participants at the 7th Interpretive Policy Analysis Conference in Tilburg, and colloquium participants at The College of New Jersey. Particular thanks are due to Melinda Roberts and to an anonymous reviewer for this journal for very generous, extensive and helpful feedback. Work on this article was supported by the University of Delaware Center for Science, Ethics and Public Policy and by Delaware EPSCoR through funds from the National Science Foundation, grant EPS-0814251.

https://doi.org/10.1017/S095382081500031X Published online by Cambridge University Press