

Quine's Population Ethics and the Need for Merely Possible People

Johan E. Gustafsson*

Draft: April 8, 2022 at 5:15 p.m.

ABSTRACT. W. V. Quine wished to restrict the interests that matter to those of actual people. Actual-Population Utilitarianism is a version of utilitarianism where, following Quine, only the interests of actual people matter. It is well known that ethical theories of this kind, which depend on what is actual, typically leads to normative variance. In this paper, I present a new objection to Actual-Population Utilitarianism. I present a case where Actual-Population Utilitarianism prescribes choices that are worse than the opposite choices for everyone whose interests matter. In addition, I present a case where total and average versions of Actual-Population Utilitarianism prescribe choices that are worse than the opposite choice for everyone (full stop). This objection also works if the Actual-Population Restriction is combined with non-utilitarian views, given some minimal assumptions.

In a brief study of population ethics, W. V. Quine professed that the only interests that matter are those of actual people. While the interests of actual future people matter, the interests of possible but non-actual future people do not: Merely possible lives don't matter.¹ This *Actual-Population Restriction* is needed in order to avoid recognizing any present yet unactual possibilities, something that Quine—for independent reasons—was eager to resist.²

* I would be grateful for any thoughts or comments on this paper, which can be sent to me at johan.eric.gustafsson@gmail.com.

¹ Quine (1978, p. 45; 1979, pp. 479–480; 1981, p. 66) wrote:

A formulation is ready to hand which sustains the moral values that favor limiting the population while still safeguarding the environment. Namely, it is a matter of respecting the future interests of people now unborn, but only of future actual people. We recognize no present unactualized possibilities.


Much the same restriction has been defended by Warren (1977, p. 285), Bigelow and Pargetter (1988, pp. 173–175), Harman (1999, p. 311), Parsons (2002, p. 142), and Cohen (2020, pp. 72–73).

² See Quine 1948, pp. 23–24.

It's well known that the Actual-Population Restriction, which depends on what is actual, can lead to *normative variance*—that is, that what *ought* to be done in a situation can depend on what *would* be done in that situation.³ This, however, is a shared problem for actualist forms of consequentialism. In this paper, I will present a new objection to the Actual-Population Restriction.

As for Quine's overall view of ethics, he favoured consequentialism, and he suggested (but did not endorse) that utilitarianism may be a systematization of our values.⁴ While we will show that my objection to the Actual-Population Restriction also works if the restriction is combined with non-utilitarian views, we start by combining the restriction with utilitarianism.

Consider the following case:



³ Arrhenius 2000, pp. 140–141 and Hare 2007, p. 503. For some objections to normative variance, see Prichard 1932, p. 26 and Carlson 1995, pp. 100–102. Hare (2007, pp. 503–508) also objects that there will be moral dilemmas if we accept the Actual-Population Restriction. But Hare's examples are neither *obligation dilemmas*—that is, situations where more than one option is obligatory—nor *prohibition dilemmas*—that is, situations where each option is wrong. (See Vallentyne 1989, p. 302.) Rather, Hare's examples are *variance dilemmas*, that is, situations where (i) there is at least one option that is not wrong, (ii) there are not two or more options each of which is obligatory and (iii) each option would be wrong if it were chosen. See Gustafsson 2019, p. 194n3.

⁴ Quine (1978, pp. 43–44; 1979, p. 478; 1981, p. 64) wrote:

There is a legitimate mixture of ethics with science that somewhat mitigates the methodological predicament of ethics. Anyone who is involved in moral issues relies on causal connections. Ethical axioms can be minimized by reducing some values causally to others; that is, by showing that some of the valued acts would already count as valuable anyway as means to ulterior ends. Utilitarianism is a notable example of such systematization.

In Bergström and Føllesdal 1994, pp. 202–204, however, Quine withheld truth from ethical statements all together.

Here, the boxes represent choice nodes. If you were to go up at node *a*, only Alice would exist and she would have a well-being of 3. If you were to go down at node *a*, there would be a second choice at node *b*. At node *b*, you have a choice between going up, which gives Alice a well-being of 2 and Bob a well-being of 5, and going down, which gives Alice a well-being of 4 and Bob a well-being of 1. The thicker arrows represent the choices that you would make at each choice node if you were to reach that node. Hence it's stipulated in this case that you would go up (even though you *could* go down) at each choice node.

If you were to reach node *b*, both Alice and Bob must be actual since they would exist in all of the then still possible outcomes. Hence the interests of both Alice and Bob would matter at node *b*. So Actual-Population Utilitarianism would prescribe going up at node *b* (since going up has a greater sum total of well-being than going down for the people whose interests would matter).

Now, note that you *would* go up at node *a* (as stipulated in the description of the case). So the only person who will actually exist is Alice. Hence she's the only person whose interests matter at node *a*. Using backward induction, we take into account that you would follow Actual-Population Utilitarianism's prescription to go up at node *b*.⁵ So we find that going up at node *a* is better than going down for everyone whose interests matter, because Alice (the only person whose interests matter at node *a*) would get a well-being of 2 if you were to go up and a well-being of just 1 if you were to go down (since you would go up at node *b*). Accordingly, Actual-Population Utilitarianism entails that you *ought to go up* at node *a*.

The trouble is that going up is worse for everyone whose interests matter (that is, Alice) than the alternative sequence of choices consisting in going down at both choice nodes. If you were to go up at node *a*, Alice would get a well-being of 3, but, if you were to go down at both choice nodes, Alice would get a well-being of 4. Thus the choices that Actual-Population Utilitarianism prescribes in this case (going up at each choice

⁵ For backward induction, see Selten 1975 and Rosenthal 1981, pp. 94–95. Both cases in this paper are *BI-terminating* decision problems—that is, the choices that are prescribed by backward induction are final in the sense that they are not followed by any further choices. See Rabinowicz 1998, p. 101. Crucially, in BI-terminating decision problems, the choices that are prescribed by backward induction can be defended with very minimal assumptions. Notably, we don't need the controversial assumption that agents would choose rationally at nodes that can only be reached through irrational choices; see Broome and Rabinowicz 1999, pp. 240–241.

node) are worse than the opposite choices (going down at each choice node) for everyone whose interests matter (everyone who actually exists). We have a violation of the following principle:

Weak Sequential Status-Confined Pareto If (i) outcome X better than outcome Y for everyone whose interests matters and (ii) Y is the outcome of an available sequence of choices, then it is not the case that an option with outcome X ought to be chosen.

Violations of this principle are worrying, since they entail that the prescriptions of the violating theory make things worse for the only people whose interests matter. The choices that Actual-Population Utilitarianism prescribes in Case One are worse than the opposite choices for everyone whose interests matters according to the Actual-Population Restriction. This can't be right.⁶

So far, we've assumed that the Actual-Population Restriction would be combined with utilitarianism. But, given the restriction to the interests of actual people, the objection to the Actual-Population Restriction need only fairly minimal ethical assumptions.

To reach the conclusion that the outcome of going up is better than the outcome of going down at node b , we only need the following principle:⁷

Weak Anonymous Status-Confined Pareto If (i) outcome X is better than outcome Y for everyone whose interests matter and (ii) Y is just like outcome Z except that the identities of some people whose interests matter have been permuted, then X is better than Z .

Let $\langle u, v \rangle$ denote an outcome where Alice gets a well-being of u and Bob gets a well-being of v . Given that the interests of both Alice and Bob matter (since they would both be actual at node b), we find that $\langle 2, 5 \rangle$ is like $\langle 5, 2 \rangle$ except that the identities of some people whose interests matter have been permuted. Since $\langle 5, 2 \rangle$ is better than $\langle 4, 1 \rangle$ for everyone whose interests would matter at node b , Weak Anonymous Status-Confined Pareto entails that $\langle 2, 5 \rangle$ is better than $\langle 4, 1 \rangle$. Accordingly, the

⁶ Don't get distracted by the observation that, if you were to go down at each node, then both Alice and Bob would actually exist and the interests of both of them would matter. The crucial thing for Actual-Population Utilitarianism is that you actually won't go down at node a and, therefore, Bob doesn't actually exist and his merely possible interests don't matter.

⁷ This principle is variation of a principle in Sen 1970, p. 153.

outcome of going up at node b , $\langle 2, 5 \rangle$, is better than the outcome of going down at that node, $\langle 4, 1 \rangle$.

Similarly, to reach the conclusion that the outcome of going up at node a is better than the outcome of going down at that node, we only need Weak Anonymous Status-Confined Pareto and backward induction. Since you would go up at node b , we find by backward induction that Alice would get a well-being of 2 if you were to go down at node a , which is lower than her well-being would be if you were to go up at node a . Since Alice is the only person whose interests matter (she's the only one who actually exists), we have, by Weak Anonymous Status-Confined Pareto, that the outcome of going up at node a is better than the outcome of going down at that node.

Hence the above objection works against the Actual-Population Restriction in combination with any consequentialist theory that satisfies Weak Anonymous Status-Confined Pareto.

So far, we haven't relied on any specific version of utilitarianism. But, if we do so, we can strengthen the objection. Given either a total or an average version of Actual-Population Utilitarianism, it violates the following principle:

Weak Sequential Fixed-Population Pareto If (i) outcome X better than outcome Y for everyone who exists in these outcomes, (ii) the same people exist in X and Y , and (iii) Y is the outcome of an available sequence of choices, then it is not the case that an option whose outcome is X ought to be chosen.


Violations of this principle should be even more worrying than the violation of Weak Sequential Status-Confined Pareto in Case One. If you go up at node a of Case One, the only person whose interests matter (Alice) is worse off than if you had gone down at all choice nodes. But, if you had gone down at all choice nodes and realized the dominating outcome, there would have been an additional person (Bob) whose interests would have mattered. On the other hand, if Weak Sequential Fixed-Population Pareto is violated, the dominating outcome has the same population as the dominated outcome.

To see how we get violations of Weak Sequential Fixed-Population Pareto, we start with the total version of Actual-Population Utilitarianism:

Total Actual-Population Utilitarianism An outcome X is at least as good as an outcome Y if and only if the sum total of well-being in X for people who actually exist and who also exist in X is at least as great as the sum total of well-being in Y for people who actually exist and who also exist in Y .

In other words, this view is the same as standard total utilitarianism except that the well-being of people who don't belong to the actual population are ignored.

Consider the following case:



At node b , it's stipulated that you would go up if you were to reach that node. So, if you were to reach node b , the actual population would include both Alice and Bob. Accordingly, Total Actual-Population Utilitarianism would prescribe going up at node b , since the total well-being for Alice and Bob is 8 if you go up but only 3 if you go down.

At node a , it's also stipulated that you would go up if you were to reach that node. So the actual population only includes Alice. Using backward induction, we take into account that, if you were to reach node b , you would (following Total Actual-Population Utilitarianism) go up at that node. Alice gets a well-being of 2 if you go up at node a , and she would get a well-being of 1 if you were to go down at node a (since you would go up at node b). Accordingly, Total Actual-Population Utilitarianism prescribes going up at node a , since Alice is the only person in the actual population. But then we have a violation Weak Anonymous Status-Confined Pareto, since we end up with an outcome where only Alice exists and where her well-being is 2 but, if you had gone down at all choice nodes, only Alice would exist and her well-being would have been 3.

Next, we turn to the average version of Actual-Population Utilitarianism:

Average Actual-Population Utilitarianism An outcome X is at least as good as an outcome Y if and only if the average of well-being in X for people who actually exist and who also exist in X is at least as great as the average of well-being in Y for people who actually exist and who also exist in Y .

In other words, this view is the same as standard average utilitarianism except that people who don't belong to the actual population don't count towards the average of well-being.

Consider, once more, Case Two. At node b , since you would go up at that node, the actual population would include both Alice and Bob. So the average well-being in the outcome of going up for the actual population is 4. And the average well-being in the outcome of going down for those in the actual population who also exist in that outcome (namely, just Alice) is 3. (Bob would be actual if you were to reach node b , but he does not exist in the outcome of going down.) Accordingly, Average Actual-Population Utilitarianism prescribes going up at node b .

At node a , since you would go up at that node, the actual population only includes Alice. Using backward induction, we take into account that, if you were to reach node b , you would (following Average Actual-Population Utilitarianism) go up at that node. So the average well-being in the outcome of going up for the actual population is 2. And the average well-being in the outcome of going down for those in the actual population who also exist in that outcome (namely, Alice) is 1. (Bob would exist if you were to go down at node a , but he's not actual.) Accordingly, Average Actual-Population Utilitarianism prescribes going up at node a .

We find that Average Actual-Population Utilitarianism prescribes the same options in Case Two as Total Actual-Population Utilitarianism. So, like Total Actual-Population Utilitarianism, Average Actual-Population Utilitarianism violates Weak Anonymous Status-Confined Pareto.

Hence, even though Quine's Actual-Population Restriction may be advantageous for our ontology, it leaves us with an implausible population ethics. The ban on merely possible people comes with a significant cost to actual people.

I wish to thank Hilary Greaves for valuable comments.

References

- Arrhenius, Gustaf (2000) *Future Generations: A Challenge for Moral Theory*, Ph.D. thesis, Uppsala University.
- Bergström, Lars and Dagfinn Føllesdal (1994) 'Interview with Willard Van Orman Quine in November 1993', *Theoria* 60 (3): 193–206.
- Bigelow, John and Robert Pargetter (1988) 'Morality, Potential Persons and Abortion', *American Philosophical Quarterly* 25 (2): 173–181.
- Broome, John and Wlodek Rabinowicz (1999) 'Backwards Induction in the Centipede Game', *Analysis* 59 (4): 237–242.
- Carlson, Erik (1995) *Consequentialism Reconsidered*, Dordrecht: Kluwer.
- Cohen, Daniel (2020) 'An Actualist Explanation of the Procreation Asymmetry', *Utilitas* 32 (1): 70–89.
- Gustafsson, Johan E. (2019) 'Is Objective Act Consequentialism Satisfiable?', *Analysis* 79 (2): 193–202.
- Hare, Caspar (2007) 'Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?', *Ethics* 117 (3): 498–523.
- Harman, Elizabeth (1999) 'Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion', *Philosophy & Public Affairs* 28 (4): 310–324.
- Parsons, Josh (2002) 'Axiological Actualism', *Australasian Journal of Philosophy* 80 (2): 137–147.
- Prichard, H. A. (1932) *Duty and Ignorance of Fact*, London: Humphrey Milford.
- Quine, W. V. (1948) 'On What There Is', *The Review of Metaphysics* 2 (5): 21–38.
- (1978) 'On the Nature of Moral Values', in Alvin I. Goldman and Jaegwon Kim, eds., *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*, pp. 37–45, Dordrecht: Reidel.
- (1979) 'On the Nature of Moral Values', *Critical Inquiry* 5 (3): 471–480.
- (1981) 'On the Nature of Moral Values', in *Theories and Things*, pp. 55–66, Cambridge, MA: Harvard University Press.
- Rabinowicz, Wlodek (1998) 'Grappling with the Centipede: Defence of Backward Induction for BI-Terminating Games', *Economics and Philosophy* 14 (1): 95–126.
- Rosenthal, Robert W. (1981) 'Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox', *Journal of Economic Theory* 25

(1): 92–100.

Selten, R. (1975) 'Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,' *International Journal of Game Theory* 4 (1): 25–55.

Sen, Amartya K. (1970) *Collective Choice and Social Welfare*, San Francisco: Holden-Day.

Vallentyne, Peter (1989) 'Two Types of Moral Dilemmas,' *Erkenntnis* 30 (3): 301–318.

Warren, Mary Anne (1977) 'Do Potential People Have Moral Rights?,' *Canadian Journal of Philosophy* 7 (2): 275–289.