## *The Need for Merely Possible People*<sup>\*</sup>

## Johan E. Gustafsson<sup>†</sup>

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 lead to normative variance. In this paper, I put forward a new objection to Actual-Population Utilitarianism. I present a case in which Actual-Population Utilitarianism prescribes choices that are worse for everyone whose interests matter than the opposite choices. Moreover, I put forward a case in which total and average versions of Actual-Population Utilitarianism prescribe choices that are worse than the opposite choice for everyone (full stop). This objection also applies if the Actual-Population Restriction is combined with non-utilitarian views, given some minimal assumptions.

In his brief study of population ethics, W. V. Quine declared that the only interests that matter are those of actual people. While the interests of future people matter, the interests of possible yet non-actual future people do not. That is, merely possible people do not matter.<sup>1</sup> 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.<sup>2</sup>

\* Forthcoming in Journal of Ethics and Social Philosophy.

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

<sup>1</sup> Quine ("On the Nature of Moral Values," 45) writes:

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 is defended by Warren ("Do Potential People Have Moral Rights?," 285), Bigelow and Pargetter ("Morality, Potential Persons and Abortion," 173–5), Harman ("Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion," 311), Parsons ("Axiological Actualism," 142), and Cohen ("An Actualist Explanation of the Procreation Asymmetry," 72–3).

<sup>2</sup> See Quine, "On What There Is," 23–4.

It is 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.<sup>3</sup> This, however, is a shared problem for actualist forms of consequentialism. In this paper, I will present a new problem for the Actual-Population Restriction.

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

Consider the following case:

<sup>3</sup> Arrhenius, *Future Generations: A Challenge for Moral Theory*, 140–1 and Hare, "Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?," 503. For some objections to normative variance, see Prichard, *Duty and Ignorance of Fact*, 26 and Carlson, *Consequentialism Reconsidered*, 100-2. Hare ("Voices," 503–8) also objects that there will be moral dilemmas if we accept the Actual- Population Restriction — for instance, a choice between (i) creating Alice at a negative level of well-being and (ii) creating Bob at a negative level of well-being. But Hare's ex amples 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, "Two Types of Moral Dilemmas," 302.) Rather, Hare's exam ples 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, "Is Objective Act Consequentialism Satisfiable?," 194n3.)

<sup>4</sup> Quine ("On the Nature of Moral Values," 43–4) writes:

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, "Interview with Willard Van Orman Quine in November 1993," 202–4, however, Quine withholds truth from ethical statements altogether.



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 a, you do not (at the time you face node a) have any voluntary control of what you would choose later 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 wellbeing of 1. The thicker arrows represent the choices you would make at each choice node if you were to reach that node. Hence, in this case, it is stipulated that you would go up (even though you *could* go down) at each choice node. Moreover, you know this in advance — that is, you know at node a that you would go up at node b.

The sequential form of Case One is crucial.<sup>5</sup> The choice at node *a* determines whether Bob will exist. At node *b*, Bob already exists (or, at least, his existence is guaranteed). Accordingly, Bob can only be created in one way (that is, by going down at node *a*).

If you were to reach node b, both Alice and Bob must then 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.<sup>6</sup> Accordingly, Actual-Population Utilitarianism would prescribe going up at node b

<sup>5</sup> If this were a synchronic choice between the three potential outcomes, we would merely have a case of normative variance — that is, we would find that, if the top outcome is chosen, only the bottom outcome would be permitted but, if either of the two lower outcomes were chosen, only the middle outcome would be permitted. This may be weird, but it's not a violation of Weak Sequential (or Weak Anonymous) Status-Confined Pareto.

<sup>6</sup> We are assuming here that, if you were to reach node *b*, the people that would then be actual (and, thus, whose interest would matter) would be the people who would exist if you were to do what you would do at that node. This is consistent with Hare's ("Voices from Another World: Must We Respect the Interests of People Who Do Not, and Will Never, Exist?," 503) *Strong Actualism* (roughly, the view that you should assess all options at a choice node taking into account only the people who would exist if you did what you would do at that node), since this is a separate application of the theory at a new (since going up has a greater sum total of well-being than going down for the people whose interests would matter).

Note that you *will* 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 is the only person whose interests matter at node *a*. Since you know at node *a* what you would do if you were to reach node *b*, you can use *backward induction*, which is to predict what you would choose at later choice nodes and to take those predictions into account when you choose at earlier nodes.<sup>7</sup> To use backward induction to determine what you ought to do is a form of *actualism*, since it makes what you ought to do now depend on what (you predict) you would do in the future.<sup>8</sup> Actualism is a rival to *possibilism* — the view that what you ought to do now depends on what you could (rather than just what you would) do in the future.<sup>9</sup> (To avoid terminological confusion, note that actualism, in the sense of taking into

choice node. Understood in this way, the actualism we consider in this paper is Strong Actualism. Hare's ("Voices," 502) *Weak Actualism* (roughly, the view that you should assess each option at a choice node taking into account only the people who would exist if you chose that option) leads to violations of some compelling principles of deontic logic (see Hare "Voices," 504–6). An alternative way of applying the restriction (call it *Super-Strong Actualism*) is to apply it rigidly to the actual world even at choice nodes that won't be reached — for instance, to apply it relative to the world where you go up at node *a* even at node *b*. But how could this be how you would apply the theory at node *b*? How would you know that you are not in the actual world? What is relevant for reasoning by backward induction is how the theory would be applied at each node. Furthermore, applying the restrictions this way may lead to Pareto violations at the choice nodes that would not be reached. To see this, suppose that Alice would, instead of 4, get a well-being of 3 if you went down at node *b*. Then, going down at node *b* would not affect those whose interests matter (that is, Alice), but Bob will be worse off than if you went up at that node.

<sup>7</sup> For backward induction, see von Neumann and Morgenstern, *Theory of Games and Economic Behavior*, 116–17, Selten, "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," and Rosenthal, "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox," 94–5. The first two cases in this paper are *BI-terminating* decision problems — that is, the choices that are prescribed by backward induction in these cases are final in the sense that they are not followed by any further choices. (See Rabinowicz, "Grappling with the Centipede: Defence of Backward Induction for BI-Terminating Games," 101.) Crucially, in BI-terminating decision problems, the choices that are prescribed by backward induction can be defended with very minimal assumptions. Notably, we do not need the controversial assumption that agents would choose rationally at nodes that can only be reached through irrational choices. See Broome and Rabinowicz, "Backwards Induction in the Centipede Game," 240–1.

<sup>8</sup> Jackson and Pargetter, "Oughts, Options, and Actualism," 233.

<sup>9</sup> Jackson and Pargetter, "Oughts, Options, and Actualism," 233.

account what you would do in the future, is distinct from the Actual-Population Restriction — that only the interests of actual people matter.) Actualism fits better with the motivation for the Actual-Population Restriction than possibilism, since actualism restricts the morally relevant acts to those that are actual (or would be actual). So we assume actualism for now, but, later on, we'll explore the Actual-Population Restriction given possibilism.

Using backward induction, you take into account that you would follow Actual-Population Utilitarianism's prescription to go up at node b. So you 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 3 if you were to go up and a well-being of just 2 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 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 is better than outcome Y for everyone whose interests matter and (ii) X is the outcome of an available sequence of choices, then it is not the case that each choice in a sequence of choices with outcome Y ought to be made.

Violations of this principle are worrying, since they entail that, for the only people whose interests matter, the prescriptions of the violating theory would make things worse. The choices that Actual-Population Utilitarianism prescribes in Case One are worse than the opposite choices for everyone whose interests matter according to the Actual-Population Restriction. This can't be right.<sup>10</sup>

<sup>10</sup> Don't be distracted by the observation that, if you were to go down at each node,

So far, we have assumed that the Actual-Population Restriction would be combined with utilitarianism. But the objection to the Actual-Population Restriction needs 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 need only the following principle:<sup>11</sup>

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 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 need only backward induction and Weak Anonymous Status-Confined Pareto. 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 is the only one who actually exists), we find, 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 backward induction and any consequentialist theory that satisfies Weak Anonymous Status-Confined Pareto. (As we

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 will not go down at node *a* and, therefore, Bob does not actually exist and his interests do not matter.

<sup>&</sup>lt;sup>11</sup> This principle is a variation of a principle in Sen, *Collective Choice and Social Welfare*, 153.

shall see later on, the objection also works against possibilist consequentialist theories.)

So far, we have not relied on any specific version of utilitarianism. But, if we do, 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* is 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 each choice in a sequence of choices whose outcome is *X* ought to be made.

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 and thereby violate Weak Sequential Status-Confined Pareto, 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 do not belong to the actual population is ignored.

Consider the following case:



At node a, you do not (at the time you face node a) have any voluntary control of what you would choose later at node b. At node b, it is stipulated that you would go up if you were to reach that node. And you know this in advance — that is, you know at node a that you would go up at node b. 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 is likewise stipulated that you will go up at that node. So the actual population includes only Alice. Using backward induction, you 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 of Weak Sequential Fixed-Population 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 do not belong to the actual population do not count toward 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 will go up at that node, the actual population includes only Alice. Using backward induction, you 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 is 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 Sequential Fixed-Population Pareto.

So far, we have relied on backward induction and actualism. It may be objected that the Actual-Population Restriction and Actual-Population Utilitarianism avoid trouble in Cases One and Two if they are instead coupled with possibilism. Given possibilism, you take into account all the things you could do and choose according to one of the optimal plans you could possibly follow. That is, (i) you consider the outcomes of all available plans and assess which of these outcomes is optimal in a choice between all of them and (ii) you ought to choose in accordance with a plan whose outcome is optimal — without taking into consideration whether you would later depart from that plan.<sup>12</sup>

Given this form of possibilism, it is no longer the case that you ought to go up at node a in Case One. Does this block the earlier objection to the Actual-Population Restriction? To see that it does not, consider once more Case One — but now we mark what ought to be done given

<sup>&</sup>lt;sup>12</sup> In decision theory, this approach is known as *naive choice*. See Pollak, "Consistent Planning," 202–3 and Hammond, "Changing Tastes and Coherent Dynamic Choice," 162.

possibilism with dashed lines (the thick lines still denote what you would do at each choice node):



Since you will actually go up at node *a*, only Alice is actual. So only Alice's interests matter at node a. The best you can do for Alice is to go down at each choice node. At node a, Weak Anonymous Status-Confined Pareto entails that the outcome of going down at each choice node is better than the outcome of any other available sequence of choices. So, given possibilism, you ought to go down at node a. And, if you were to reach node *b*, you would go up at that node. So both Alice and Bob would be actual, and their interests would matter if you were to reach node b. So, at node b, it follows (in the same way as before) by Weak Anonymous Status-Confined Pareto that the outcome of going up is better than the outcome of going down. So, given possibilism, you ought to go up at node b. Hence each choice in the sequence of choices consisting in going down at node *a* and going up at node *b* ought to be made given possibilism. But the outcome of going up at node a is better for everyone whose interests matter (namely, Alice) than the outcome of going down at node *a* and up at node *b*. Hence we still have a violation of Weak Sequential Status-Confined Pareto in Case One.

But how about the objection that Total and Average Actual-Population Utilitarianism both violate Weak Sequential Fixed-Population Pareto in Case Two? Actually, given possibilism, that objection no longer works in Case Two. To see this, consider that case once more but with what ought to be done given possibilism marked with dashed lines:



At node *a*, you will go up. So, at node *a*, only Alice is actual and so only her interests matter. Then, at node *a*, Total and Average Actual-Population Utilitarianism each entails that the outcome of going down at each choice node is better than the outcome of any other available sequence of choices. So, given possibilism, you ought to go down at node *a*. But then we have no violation of Weak Sequential Fixed-Population Pareto.

Nevertheless, Total and Average Actual-Population Utilitarianism still violate Weak Sequential Fixed-Population Pareto in another case. Consider the following case:



At each choice node, you do not (at the time you face that node) have any voluntary control of what you would choose later at future nodes. It is stipulated that, at each choice node, you would go up (even though you could go down). So, at node *a*, only Alice is actual and so only her interests matter. Therefore, according to both Total and Average Actual-Population Utilitarianism, the best outcome is the outcome of going down at both node *a* and node *b*. Hence, given possibilism, you ought to go down at node *a*.

If you were to reach node b or node c, both Alice and Bob would be actual (since you would go up at those nodes). So, at nodes b and c, the best outcome of any available sequence of choices is the outcome of go-

ing up at node *b* and then down at node *c*. Therefore, according to both Total Actual-Population Utilitarianism and Average Actual-Population Utilitarianism, you ought to go up at node *b* and then down at node *c*. Hence, given possibilism, each choice in the sequence of choices consisting in going down at node *a*, going up at node *b*, and going down at node *c* ought to be made given possibilism. But the outcome of going up at node *a* is better for everyone whose interests matter (namely, Alice) than the outcome of going down at node *a*, up at node *b*, and down at node *c*. Hence, even given possibilism, we have a case where Total Actual-Population Utilitarianism and Average Actual-Population Utilitarianism violate Weak Sequential Fixed-Population Pareto.<sup>13</sup>

In summation, Quine's Actual-Population Restriction may be advantageous leaves us with an implausible population ethics. The silencing of the interests of merely possible people comes with a cost to actual people.<sup>14</sup>

<sup>13</sup> It may be objected that you actually violate the prescriptions of the possibilist versions of Total Actual-Population Utilitarianism and Average Actual-Population Utilitarianism in Case Three. But note that all available sequences of choices, in this case, would violate those prescriptions. Going up at node *a* violates the prescription at that node. Going down at node *a*, up at node *b*, and up at node *c* would violate the prescription at node *c*. Going down at node *a*, up at node *b*, and down at node *c* would violate the prescription at node *b*. Going down at node *a* and down at node *b* would violate the prescription at node *b*. Hence these possibilist theories are sequentially unsatisfiable: there is no available sequence of choices in Case Three such that, if you were to make that sequence of choices, you would not violate the theory (even though these theories are satisfiable at each choice node).

<sup>14</sup> I wish to thank John Broome, Krister Bykvist, Hilary Greaves, Caspar Hare, Wlodek Rabinowicz, Melinda Roberts, and an anonymous referee 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.
- Hammond, Peter J. (1976) 'Changing Tastes and Coherent Dynamic Choice', *The Review of Economic Studies* 43 (1): 159–173.
- 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.
- Jackson, Frank and Robert Pargetter (1986) 'Oughts, Options, and Actualism', *The Philosophical Review* 95 (2): 233–255.
- Parsons, Josh (2002) 'Axiological Actualism', *Australasian Journal of Philosophy* 80 (2): 137–147.
- Pollak, R. A. (1968) 'Consistent Planning', *The Review of Economic Studies* 35 (2): 201–208.
- 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.
- Rabinowicz, Wlodek (1998) 'Grappling with the Centipede: Defence of Backward Induction for BI-Terminating Games', *Economics and Phi*-

*losophy* 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.
- von Neumann, John and Oskar Morgenstern (1944) *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.
- Warren, Mary Anne (1977) 'Do Potential People Have Moral Rights?', *Canadian Journal of Philosophy* 7 (2): 275–289.