Probabilities in Tragic Choices

EDUARDO RIVERA-LÓPEZ
Universidad Torcuato Di Tella (Buenos Aires)

In this article I explore a kind of tragic choice that has not received due attention, one in which you have to save only one of two persons but the probability of saving is not equal (and all other things are equal). Different proposals are assessed, taking as models proposals for a much more discussed tragic choice situation: saving different numbers of persons. I hold that cases in which (only) numbers are different are structurally similar to cases in which (only) probabilities are different. After a brief defense of this claim, I conclude that some version of consequentialism seems more promising for offering a plausible solution to the probability case.

I

Two strangers, *A* and *B*, will die unless you give them one dose of a medicine. Unfortunately, you have only one dose; you can save either *A* or *B*, but not both. However, you know something else: *A* is more ill than *B*. If you give *A* the medicine, it is less probable that *A* will survive. If you give *B* the medicine, *B* will surely (or more probably) survive.

You are facing a tragic choice. Tragic choice situations are those in which every available course of action foreseeably leads to a seriously bad outcome for someone. There are many different kinds of tragic choice situations and, while some of them have been widely discussed, others have received much less attention. As far as I know, the kind of tragic choice presented in my example has received (in its pure form) no philosophical attention.[1] My aims in this article are to explore it in its pure form, to test some arguments that have been proposed for other kinds of tragic choices (specifically, situations related to saving different numbers of people), and, ultimately, to show that the consequentialist solution is the most plausible one.

In the following section, after briefly presenting some basic definitions and sketching the article's conceptual framework, I consider several arguments, taking as a model the case of saving different numbers of persons. In section III, I offer some reflections on the specificity of probabilities in tragic choices, showing that cases involving probabilities and cases involving numbers are similar in relevant ways.

Before proceeding, it may be worth making a brief comment on the practical relevance of the specific problem. There are situations in which we have to rescue people from death and in which only positive

[1] By 'pure case' I mean a tragic choice situation in which only one factor (in my example, the probability of saving) differentiates the various choices.

© 2008 Cambridge University Press
Utilitas Vol. 20, No. 3, September 2008
doi:10.1017/S0953820808003178
Printed in the United Kingdom

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

actions are involved. These situations arise in different contexts, such as accidents or natural catastrophes in which authorities must adopt one of two (or more) rescue strategies with scarce resources, as well as in situations in which doctors (or healthcare authorities) must allocate scarce organs for transplant (or some other life-saving treatment, such as dialysis) to some patients in preference to others. In these situations, several factors are usually at work simultaneously: the probabilities of rescuing, the numbers of victims, the significance of the benefit to be provided, among others. Real-life cases are impure and therefore require compromises between these factors. However, it is unlikely that we can find plausible criteria to deal with these real (impure) cases if we do not know how to deal with pure (and therefore much simpler) cases: cases in which the available options differ only with regard to one factor. In this article, I deal with the case in which that factor is the probability of saving.

II

Let us first present the case in clearer terms:

Probability Machine:

A and *B* are captured in a diabolical machine that works as follows:
If you press button 1, you give *A* a 50 percent chance of survival, but *B* dies (for sure).
If you press button 2, you save *B* (for sure), but *A* dies (for sure).
All other things are equal.[2]
You must press either button 1 or button 2, but not both.

Probability Machine is an example of what we can call 'Probability Case' (and the problem about this case we can call the 'Probability Problem'). Compare Probability Machine with the following much more discussed example:

Numbers Machine:

A, *B* and *C* are captured in a diabolical machine that works as follows:[3]
If you press button 1, you save *A*, but *B* and *C* die.
If you press button 2, you save *B* and *C*, but *A* dies.

[2] The 'all other things being equal' clause must be understood to include any other feature of the people in each group. In this case, it includes the number of persons being saved, the length of the life after being saved, the quality of life of the saved persons, and so on.

[3] This way of presenting the Numbers Case resembles Tooley's example in 'An Irrelevant Consideration: Killing Versus Letting Die', *Killing and Letting Die. Second Edition*, ed. B. Steinbock and A. Norcross (New York, 1994), p. 106. The primary difference is that each of your options is a positive action in the Numbers Case, whereas Tooley presents his example in order to discuss the symmetry (or asymmetry) between actions and omissions.

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

All other things are equal.
You must either press button 1 or button 2, but not both.

Numbers Machine is an example of what we can call 'Numbers Case'. Correspondingly, it presents the 'Numbers Problem'. I will now explore different responses to Probability Machine, taking into account some solutions that have been provided to the Numbers Problem.

Utilitarians and common-sense morality claim that you should press button 2 in Numbers Machine. This is because pressing button 2 will save more lives. Since adding different people's well-being is (at least) part of the prescribed decision-making procedure, they are aggregationists.[4]

Utilitarianism and common-sense morality are not only aggregative but also consequentialist concerning numbers, since the underlying reasoning is the following: more lives being saved is a better state of affairs than fewer lives being saved.[5] Therefore, you should press button 2, which causally leads to more lives being saved. However, not all consequentialists are aggregationists.[6] I. Hirose provides a consequentialist solution to the Numbers Problem, which he claims is not aggregative.[7] He starts from two weak assumptions: an Impartiality principle and a Pareto principle. According to the Impartiality principle, it is morally indifferent whether *A* survives and both *B* and *C* die, or whether *A* and *C* die and *B* survives (two states of affairs are equally good if they differ only in the identities of the participants). According to the Pareto principle, a state of affairs in which both *B* and *C* survive and *A* dies is better than one in which *A* and *C* die and only *B* survives (because in the first state of affairs no one is worse off and at least one person is better off). From these two premises, it follows that a state of affairs in which *B* and *C* survive and *A* dies is better than one in which *A* survives and *B* and *C* die. This argument is consequentialist because your moral duty to press button 2 follows from a conclusion about the goodness of states of affairs. However, it is not aggregative since counting or adding people is not part of your decision procedure.

[4] For a definition of 'aggregation', see I. Hirose, 'Aggregation and Numbers', *Utilitas* 16.1 (2004), p. 66 ('the combination of separate people's goods, happiness, losses, well-beings, and so on, into an objective value').

[5] Common-sense morality is not consequentialist on many moral issues. However, concerning Numbers Cases, it is clearly compatible with it.

[6] Utilitarisnism, as I understand it, is a consequentialist theory, which also is aggregative and welfarist.

[7] See I. Hirose, 'Saving the Greater Number without Combining Claims', *Analysis* 61.4 (2001), and Hirose, 'Aggregation and Numbers', pp. 68–9. This solution is admittedly inspired by what Kamm calls the 'aggregation argument' (F. M. Kamm, *Morality, Mortality*, vol. 1 (Oxford, 1993), pp. 85–7).

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

Concerning Probability Machine, utilitarianism (aggregative consequentialism) and common-sense morality will reach the same solution: you should press button 2. Utilitarians will prescribe pressing button 2 because it maximizes expected utility. Common-sense morality will reach the same solution, perhaps for less sophisticated reasons (such as having to opt for a sure choice).[8] It is also possible to apply Hirose's non-aggregative approach to Probability Machine, in the following way: (i) *A*'s receiving a 50 percent (and *B* a 0 percent) chance of survival is as good as *B*'s receiving a 50 percent (and *A* a 0 percent) chance of survival (for Impartiality); (ii) *B*'s receiving a 100 percent (and *A* a 0 percent) chance of survival is better than *B*'s receiving a 50 percent (and *A* a 0 percent) chance of survival (for Pareto); therefore, (iii) *B*'s receiving a 100 percent (and *A* a 0 percent) chance of survival is better than *A*'s receiving a 50 percent (and *B* a 0 percent) chance of survival.

Putting aside Hirose's claim that this argument is not aggregative (which is disputable), we can safely reach an initial conclusion: consequentialist arguments (including clearly aggregative, utilitarian ones) are easily transferable from Numbers Machine to Probability Machine without loss of force.

Of course many philosophers are not consequentialists, and therefore reject either the claim that we should press button 2 in Numbers Machine, or the claim that we should press button 2 *for consequentialist reasons*. John Taurek takes the most extreme position against consequentialist solutions for tragic choice cases. He holds that Numbers Machine presents us with no moral reason to press button 2, thereby saving *B*'s and *C*'s lives. On his view, there is no neutral moral perspective from which you can judge that it is morally better that two lives be saved than one. If there is no special obligation, you should give each person the highest equal chance of survival. You should therefore flip a coin.[9]

It is very plausible that Taurek would offer the same proposal about Probability Machine that he offers about Numbers Machine: that you should flip a coin. In the next section, I argue for this claim. My point here is that, concerning Probability Machine (and, therefore, the Probability Problem in general), there are at least two reasons to believe that his proposal is less appealing than in Numbers Machine. First, flipping the coin permits you to do what may not save anyone in

[8] This does not mean that utilitarianism and common-sense morality will agree in every Probability Case. It is not clear, for example, that common-sense morality would require you to press button 2 if pressing button 1 would give *A* a 98 percent chance of survival instead of a 50 percent chance.

[9] See J. M. Taurek, 'Should the Numbers Count?', *Philosophy and Public Affairs* 5.4 (1977), pp. 303–4.

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

Probability Machine (this is what will happen if you flip a coin and *A* wins, but when you press button 1 *A* does not survive). In contrast, by pressing button 2, you are sure to save one person. This is an important difference from taking Taurek's advice in Numbers Machine. In that case you are at least sure that someone will be saved. Second, imagine a modified version of Probability Machine in which *A*'s chance of survival if you press button 1 is not 50 percent but 2 percent (provided that you press button 1). It would be extremely counterintuitive to hold that you have to flip a coin in this case, since this would give you a 49 percent chance of saving nobody at all.[10]

For many people, it seems acceptable neither to be a utilitarian (or, more broadly a consequentialist) nor to have to flip a coin in Numbers Cases. Therefore, some effort has been made to find a non-consequentialist argument for saving the many in Numbers Cases. F. Kamm and T. M. Scanlon have provided a strong argument of this sort.[11] They argue as follows. First imagine a situation in which you have to choose between saving *A*'s life and saving *B*'s life (all other things being equal). In such a situation, *A* and *B* have equal claims. If you flip a coin, you are taking both claims into account. Now transform this situation into Numbers Machine. If you still believe that flipping a coin is the right procedure, you are acting as if everything were comparable to the previous situation. But it is not. Now *C*'s claim is also at stake. He will be saved if you press button 2. By flipping a coin you would be giving *C* no consideration at all. *C* would therefore reasonably reject the procedure. *C*'s legitimate claim tilts the balance in favor of button 2.

Let us examine Probability Machine in light of the Kamm/Scanlon account. We might at first think that it is consistent with the coin-flipping proposal: by flipping a coin you would be giving *A* and *B* equal consideration and, therefore, nobody could reasonably reject your decision-making procedure. However, I have already explained why I find this solution highly implausible. And in any case, on this construal the Kamm/Scanlon contractualist account yields no progress over Taurek's. A second strategy for contractualists might be to argue that you should press button 2, but not for consequentialist or aggregative reasons. This would be analogous to the contractualist solution to Numbers Machine. The analogous argument is the following. Suppose we face a situation in which everything resembles Probability Machine

[10] Admittedly, this result may not be much more counterintuitive than the idea that one should flip a coin when (in a Numbers Case) one has to choose between saving one person and saving one hundred. Still, in this case you are at least sure that one person will be saved.

[11] Kamm, *Morality, Mortality*, pp. 116–17; T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass., 1998), p. 232.

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

except that if you press button 1, *A* has a 50 percent chance of survival (and *B* dies), and, if you press button 2, *B* has a 50 percent chance of survival (and *A* dies). In this case, it is obvious that it is right to flip a coin. Now suppose you add the remaining 50 percent chance to *B*, transforming this situation into Probability Machine. If everything remained unchanged and you still believed that flipping a coin was the right procedure, *B* would have a legitimate reason to complain. She could reasonably argue that her improved chance of survival (from 50 percent to 100 percent) should have had at least some impact on the decision-making procedure. Such an increase, the argument concludes, should tilt the balance in favor of pressing button 2.

However, there is a crucial difference between Numbers Machine and Probability Machine which makes it difficult to accept this proposal. In Numbers Machine, we have *C*'s claim that he is not adequately considered in the flipping-a-coin procedure. *C* is a different person. The coin-flipping method does not take his claim into account *at all*. In contrast, in Probability Machine it is *B*'s chance of survival that rises from 50 to 100 percent (if you press button 2). When her chance increases to 100 percent, she cannot complain about not being considered by flipping a coin, because she *is* being considered. The emergence of a new person (*C*) seems to force us to change our decision-making procedure so as to account for such an emergence. But an increase in the probability of saving the same person (*B*) does not necessarily lead to the same outcome.

Some philosophers have objected that the Kamm/Scanlon argument for pressing button 2 has, against the intention of its proponents, aggregative features, but the point is controversial.[12] J. Timmermann endorses this objection and defends a distinct non-consequentialist and non-aggregative solution to the Numbers Problem. According to his method, which Timmermann calls the 'individualist lottery', you should spin a wheel of fortune that has three sections, one corresponding to each person involved: *A*, *B* and *C*. If either *B* or *C* wins, you will have to press button 2, saving both (if *B* wins, *C*'s life will be saved as a

[12] Note that, when *B*'s chances increase from 50 percent to 100 percent and you accordingly change your decision from flipping the coin to pressing button 2, you are sacrificing *A* *completely* in exchange for an increase of *B*'s chances of survival. In Numbers Machine, on the other hand, when you change from flipping the coin to pressing button 2, you are sacrificing *A* completely in exchange for giving *C* *some* consideration. All this is of course controversial. One might plausibly argue that it is not true that *C* is not taken into account by flipping the coin (in fact, you would be giving him the same chance of survival as *A* and *B*). The only reason to press button 2 has to be that *C* is *together* with *B*. This would imply that the argument is, after all, aggregative. See M. Otsuka, 'Scanlon and the Claims of the Many versus the One', *Analysis* 60.3 (2000) for this line of argument, and R. Kumar, 'Contractualism on Saving the Many', *Analysis* 61.2 (2001) and Hirose, 'Aggregation and Numbers', pp. 72–3 for the opposite view. For my purposes, I do not need to push in this direction.

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

side-effect, and, if *C* wins, *B*'s life will be saved as a side-effect). If *A* wins, you will have to press button 1, letting *B* and *C* die. This account is non-aggregative because you do not have to count how many people are saved under each alternative. You only have to give each involved person an *ex ante* equal chance of survival.[13]

Consider now how the individualist lottery could be adapted to Probability Machine. Again, the defender of this account might propose flipping a coin. The idea would be that the principle underlying the individualist lottery in Numbers Machine is that each person receives the same chance, regardless of what further conditions might help or not help such a person (in Numbers Machine, *A*'s further condition is her being alone and *B*'s and *C*'s further condition is their being together). In Probability Machine, we give each person, *A* and *B*, the same *ex ante* probability of being saved, regardless of what further conditions hold: *A*'s lower chance of survival and *B*'s higher chance of survival.

We have already seen problems with the coin-tossing solution. In addition to these problems, it might seem that this solution, in some sense, would defeat a deeper rationale of the individualist lottery – that we should give each participant the same chance of survival. By flipping a coin, we would give *B* a 50 percent chance of survival, whereas we would give *A* only a 25 percent chance. It is true that you are not responsible for the inegalitarian mechanism of the machine (which gives a 100 percent chance of survival to *B* if you press button 2, and only a 50 percent chance to *A* if you press button 1). However, it might be argued that there is no reason not to take account of the machine's mechanism when selecting the decision-making procedure. By flipping a coin, you are, in fact, *not* giving them an equal chance.

Let us assume that we want to equalize chances of survival. This will lead us to what we may call the 'equalizing lottery'. The equalizing lottery equalizes chances by incorporating the mechanism into the decision-making procedure. You run a wheel of fortune that gives a 1/3 chance of survival to *B* and a 2/3 chance to *A*. The upshot is that each then has a 1/3 chance of survival. This proposal might appear plausible, but it is vulnerable to a strong 'leveling down' objection: we obtain equality only at the expense of the overall probability of saving anyone at all. If each person has a 1/3 chance of survival, there remains a 1/3 chance that nobody survives. To make the point more obvious, suppose again that *A*'s chance of survival if you press button 1 were not 50 percent but 2 percent (*B*'s chance remaining unchanged). In this case, if you wish to apply the equalizing lottery, you should give slightly

[13] J. Timmermann, 'The Individualist Lottery: How People Count, But Not Their Numbers', *Analysis* 64.4 (2004).

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

less than a 2 percent chance to both *A* and *B*. But then you will have a slightly more than 96 percent chance of saving no one at all!

A final proposal with regard to Numbers Machine, which is not strictly consequentialist but which contains an undeniable aggregative element, is the so-called 'weighted lottery'.[14] The weighted lottery deals with Numbers Machine in the following way. Since you save two persons by pushing button 2 and only one person by pushing button 1, you should run a wheel of fortune that gives a 2/3 chance to pressing button 2 and a 1/3 chance to pressing button 1. The general idea is that the weight of each alternative should be proportional to the number of persons saved. You would be giving some consideration to *A* (which is denied by the utilitarian solution); but, at the same time, you would be giving (proportionally) more consideration to *B* and *C*, because (and only because) of their number.[15]

Let us then see how we might apply the idea of the weighted lottery to Probability Machine. As in Numbers Machine, we start from the premise that pressing button 2 is more valuable than pressing button 1. In Numbers Machine this is so because you save more persons by pressing button 2. In Probability Machine this is so because it is more likely that you save a human life by pressing button 2. But you do not want to give *A* no moral consideration, and therefore you give her a weight that is proportional to how less good it would be to press button 1 than to press button 2. With Probability Machine, this idea might be applied as follows. If pressing button 1 gives *A* a 50 percent chance of survival and pressing button 2 gives *B* a 100 percent chance, then you should give a 2/3 chance to pressing button 2 and a 1/3 chance to pressing button 1. It is as if pressing button 2 were twice as valuable as pressing button 1 because *B*'s chance of survival is twice *A*'s chance. This solution would give a 2/3 chance of survival to *B* and a 1/6 chance of survival to *A*.

It must be noted that the weighted lottery inherits one of the problems of the equalizing lottery: the possibility of saving no one at all. However, the chance of not saving anyone is lower on this account: if by your pressing button 1 *A*'s chance of survival is 50 percent, the overall chance of not saving anyone under the weighted lottery is 16.66 percent (against 33.33 percent under the equalizing lottery). If *A*'s chance were 2 percent, the probability of not saving anyone would be 1.92 percent (against 96 percent under the equalizing lottery). I am not sure

[14] See J. Broome, 'Selecting People Randomly', *Ethics* 95.1 (1984), p. 55. The idea is also discussed in Kamm, *Morality, Mortality. Volume 1,* pp. 128–9 and Scanlon, *What We Owe to Each Other*, pp. 233–4.

[15] As Timmermann rightly points out, this solution is pragmatically identical to the individualist lottery. They are, however, theoretically different because the individualist lottery is (or at least attempts to be) non-aggregative.

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

about the ultimate plausibility of this solution. Still, it is worth noting that, even if this approach is not strictly consequentialist, it is not strictly deontological either. The relative goodness of states of affairs produced by each of the possible alternatives of action is still a relevant consideration of the decision-making procedure. It is just that this consideration is not the only one: a non-consequentialist consideration constrains your decision by requiring you to give each person a proportional chance. Rather than being regarded as deontological, such an account might better be thought of as a restricted form of consequentialism.

Our conclusions so far are these. Consequentialist and common-sense accounts work both with Numbers Machine and with Probability Machine. Not so with non-consequentialist arguments. They might work with numbers, but they are substantially less persuasive with probabilities.

III

We might think that the difficulty of finding a non-Taurekian but, at the same time, non-consequentialist solution to the Probability Problem is not due to a problem with non-consequentialism, but rather due to the assumption that Probability Cases are relevantly analogous to Numbers Cases. Translating arguments from Numbers to Probability Cases might be a mistake. Perhaps there is a non-consequentialist argument showing that we should press button 2 in Probability Machine – an argument that is not derived or adapted from some other argument concerning Numbers Machine. Such an argument might exist. Still, there are a number of reasons that suggest that Numbers and Probability Cases are, at a fundamental level, sufficiently similar to warrant the assumption that we should take solutions to Numbers Machine as models for dealing with Probability Machine.

The first reason is that the underlying utilitarian argument for pressing button 2 in Numbers Machine and in Probability Machine is the same. It is not that you should press button 2 in Numbers Machine just because you should be willing to aggregate interpersonal utility. You should do so because you should be willing to aggregate expected utility in general, and, therefore, interpersonal utility as well. Hence, you should also press button 2 in Probability Machine.

The same holds in the opposite view. In my view, the fundamental idea behind Taurek's refusal to count people in numbers cases has nothing to do with adding up different persons into a whole. It has to do with the impossibility of assuming a neutral perspective from which to compare the well-being of different persons. According to Taurek, there are only personal, agent-relative perspectives. The thesis that

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

numbers do not count is a corollary of this more general view. This is why Taurek does not require us to choose to save *C*'s life over saving *B*'s arm.[16] I cannot compare *C*'s loss *vis-à-vis B*'s loss from a neutral point of view. For *B*, her losing an arm is (or at least can be) worse than *C*'s death. The same reasons that prevent us from adding persons in Numbers Machine also work in Probability Machine and prevent us from comparing the harms or benefits that will come to *A* or *B*.

Moreover, any numbers skeptic who is willing to aggregate utility in cases in which the number of persons is the same (as in Probability Machine) is likely to fall into normative inconsistencies. Suppose you have to choose between

(i) giving *A* no chance of survival, and *B* and *C* each a 50 percent chance,
(ii) giving *A* a 35 percent chance of survival, *B* no chance and *C* a 50 percent chance,

or

(iii) giving *A* a 70 percent chance of survival, and *B* and *C* no chance.

This is a simplified version of Michael Otsuka's argument against numbers skepticism, which I have adapted to probabilities.[17] If you are a numbers skeptic but, at the same time, believe that you can make pairwise comparisons of probabilities of survival (you press button 2 in Probability Machine), then you will prefer (i) to (ii) (it is better to give *B* a 50 percent chance than to give *A* a 35 percent chance). You will also prefer (ii) to (iii) (*C*'s loss of falling from 50 to 0 percent is worse than *A*'s gain from 35 to 70 percent). However, you will also have to prefer (iii) to (i), because giving *A* a 70 percent chance of survival is better than giving *B* or *C* a 50 percent chance (and you cannot add the percentages of each). You are caught in a circle of intransitive (moral) preferences. From this argument Otsuka concludes that numbers skepticism has a flaw and that we should reject it. My conclusion is that the problem is one of consistency between being a numbers skeptic and, at the same time, accepting pairwise comparisons of harm (in terms of probabilities of suffering the same harm or in terms of the amount of harm).

All this strongly suggests that there is a close connection between how we deal with Numbers Machine and Probability Machine.

[16] See Taurek, 'Should the Numbers Count?', p. 302.
[17] See M. Otsuka, 'Skepticism about Saving the Greater Number', *Philosophy & Public Affairs* 32.4 (2004), pp. 421–3. Otsuka's example does not deal with different probabilities but with different amounts of harm, concretely, different numbers of limbs that can be restored.

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press

Therefore, proposals regarding the first case can work as models for exploring proposals regarding the second.

CONCLUSION

As I have tried to show in the previous section, the Numbers and Probability Problems are in relevant ways similar. In arguments for and against counting persons in Numbers Cases, the central issue is whether different people's well-being is commensurable or not. And this same issue is central in cases in which numbers are not what makes the difference between one's options (such as the Probability Problem). On the basis of these claims, my claim has been that, if we reject consequentialism and aggregationism but, at the same time, also reject Taurek's conclusions about Numbers Machine, we may well reach a plausible solution to the Numbers Problem. But we will not be able to extend such a solution to the Probability Problem. This problem remains intractable for deontological theories. Some version of consequentialism therefore seems more promising.[18]

erivera@utdt.edu

[18] I presented an earlier version of this article as a paper at the symposium 'Current Problems in Moral and Legal Theory', Universidad Torcuato Di Tella (June 2007). I am grateful to the audience for many helpful remarks. I also want to thank Marcelo Ferrante, Joshua Gert, Nora Muler, Michael Otsuka, the editor and an anonymous reviewer for their comments and advice.

https://doi.org/10.1017/S0953820808003178 Published online by Cambridge University Press