

THE EX ANTE PARETO PRINCIPLE*

I introduce the idea of pareto superiority by comparing two scenarios.

Two rabies. You have two patients before you, Alan and Bob, who both have rabies. You have only one dose of medicine, which you can give to either Alan or Bob. The medicine will cure the rabies, but has some mild unwelcome side-effects, such as temporary nausea and tiredness. In patients susceptible to eczema, the medicine also causes an itchy skin rash which always clears up within a couple of days. Alan and Bob's situations are relevantly similar, except that Bob is susceptible to eczema so if he takes the medicine then he will get the skin rash as well as the other side-effects. Below is a table rating the outcomes of your different actions for Alan and Bob – who we can assume are the only people affected by your choice of action in any significant way. You can take the numbers as measures of well-being: exactly how well-being should be measured is of course controversial, but my argument is independent of this debate. Take the numbers to represent well-being as measured in your preferred way:

	<i>Outcome for Alan</i>	<i>Outcome for Bob</i>
Give the medicine to Alan	0.9	0
Give the medicine to Bob	0	0.85

Table 1

* Many thanks to the Leverhulme Trust, who funded this research. I am also very grateful for many helpful comments and questions from audiences at Belfast, the James Allen's Girls' School, the LSE, Munich (MCMP), and Nottingham. In particular, thanks to Nicholas Baigent, Luc Bovens, Richard Bradley, Seamus Bradley, Chloé de Canson, Goreti Faria, Caspar Hare, Alexander Kaiserman, Ofra Magidor and Mike Otsuka. Finally thanks to an anonymous reviewer for this journal.

What should you do?

One rabies, One Healthy. You are a doctor with two patients before you, Alan and Carol. Alan has rabies, and Carol is healthy. You have only one dose of medicine, which you know will cure Alan (though it will leave him with some side-effects). Carol is perfectly healthy without the medicine, and giving her the medicine will simply give her the side-effects. Here are the outcomes of your different actions for the people involved:

	<i>Outcome for Alan</i>	<i>Outcome for Carol</i>
Give the medicine to Alan	0.9	1
Give the medicine to Carol	0	0.9

Table 2

What should you do?

The choice in the first scenario is difficult. Giving the medicine to Alan has a better outcome for Alan, but giving the medicine to Bob has a better outcome for Bob. Which outcome is better overall? Is there some way of aggregating the well-being of individuals across an outcome, so that we can assess the outcome's overall goodness? Some – certain utilitarians – would recommend simply summing the well-being across the individuals in any given outcome, without worrying about how that well-being is distributed. Others – the priortarians – place particular weight on the well-being of the worst off. And still others – the egalitarians – pay attention to how equally the total well-being is distributed amongst the people involved. Which of these approaches – if any – is right, is a controversial matter.

The choice in the second scenario, on the other hand, is a no-brainer. Giving the medicine to Alan (rather than to Carol) is better for Alan, and it is also better for Carol. It is better for everyone concerned. Economists would classify the outcome in which Alan gets the medicine as 'pareto superior' to the outcome in which Carol gets the medicine. Here is a definition of pareto superiority (for outcomes): an outcome O_1 is pareto superior to another outcome O_2 iff every individual is at least as well off in O_1 as in O_2 , and furthermore at least one individual is better off in O_1 than in O_2 . When ranking outcomes for the purposes of welfare economics, it is generally taken for granted that whenever one outcome O_1 is pareto superior to another outcome O_2 , O_1 is better than O_2 .¹ And where you have a choice between an action that you are certain will produce an outcome O_2 and an action that you are certain will produce a better outcome O_1 , of course you should carry out the action that will produce the better outcome.

In real-life cases, the certainty is usually missing: you cannot be certain what outcomes your actions will have. With this in mind, let's consider a more realistic version of *One Rabies, One Healthy*. In a realistic scenario, you won't be *certain* that Alan has rabies: even if he has all the symptoms, and the test results are positive, there is the possibility that some other disease has caused similar symptoms, and that the testing process went wrong somewhere. Then taking the medicine might make Alan's situation worse: his disease will not be cured, and he will suffer from the side-effects of the medicine. Similarly, there is the possibility that Carol is not healthy, but merely seems it: perhaps she too has rabies but is not exhibiting any symptoms, in which case the rabies medicine is just what she needs. Furthermore, you are unlikely to be *certain* that

¹ "...virtually all [welfarists] agree that this ranking should satisfy the principle of Pareto-superiority (strong Pareto)." Matthew D. Adler, *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis* (Oxford: OUP, 2011), p.52.

the medicine will have mild side-effects. Perhaps if Alan takes the medicine, it will have no side-effects at all – or perhaps it will provoke a violent and fatal allergic reaction.

One way (pioneered by Leonard Savage² (1954)) to represent your uncertainty about the outcomes of your actions, is to have different columns representing different possible states of the world. Below I give a table along these lines, with just two possible states of the world: the state in which Alan has an extreme allergy to the medicine, such that the medicine would cause sudden death, and the state in which Alan does not have such an allergy, in which case (let's assume) the side-effects will be mild. Of course the table is still unrealistically simple – but it serves to make my point:

	<i>State S₁: Alan is not allergic to the medication</i> <i>Pr(S₁) = 0.999</i>		<i>State S₂: Alan is allergic to the medication</i> <i>Pr(S₂) = 0.001</i>	
	Outcome for Alan	Outcome for Carol	Outcome for Alan	Outcome for Carol
Give the medicine to Alan	0.9	1	-0.2	1
Give the medicine to Carol	0	0.9	0	0.9

Table 3

We can see from the table that in state S₂ (where Alan is allergic to the medicine), Alan has a worse outcome if you give him the medicine than if you give it to Carol. Of course, in state S₂ Alan is going to die whether he gets the medicine or not, but we can suppose that it would be nicer for him to have a few extra hours of life than for him to die rapidly and unexpectedly in your surgery. In state S₂, the outcome of giving the medicine to Alan is not pareto superior to

² Leonard J. Savage, *The Foundations of Statistics* (New York: Dover, 1972).

the outcome of giving the medicine to Carol. And you cannot be certain – at the time when you are deciding what to do – whether state S_1 or state S_2 obtains. Thus you cannot be certain that giving the medicine to Alan will result in an outcome that is pareto superior to the outcome that would have obtained had you instead given the medicine to Carol. Does the principle of pareto superiority then not offer any guidance in this scenario?

There is a natural way to extend the principle so that it does offer guidance in this scenario. Rather than directly comparing the outcomes for the individuals under the different actions, we compare the *expected* outcomes for the individuals under the different actions. The expected outcome for an individual under an action is the sum of – for each possible state of the world – the probability of that state's obtaining multiplied by the outcome for the individual at that state under the relevant action.³ The table below gives the expected outcomes for the scenario we are considering:

	<i>Expected Outcome for Alan</i>	<i>Expected Outcome for Carol</i>
Give the medicine to Alan	$(0.999)*(0.9) + (0.001)*(-0.2)$ $= 0.8991 - 0.0002 = 0.8989.$	1
Give the medicine to Carol	0	0.9

Table 4

Now we classify an *action* A_1 (rather than an outcome) as pareto superior to another action A_2 iff for every person the expected outcome under action A_1 is at least as good as the expected outcome under A_2 , and for at least one person, the expected outcome under action A_1 is better

³ Complications arise in cases where performing a given action changes the likelihood of a state's obtaining: we will restrict our attention to cases where the states and actions are independent.

than the expected outcome under A_2 . Thus we can see that from table 4 that giving the medicine to Alan is pareto superior to giving the medicine to Carol.

This concept – of pareto superiority as a relation between *actions* – is the focus of this paper. The concept is sometimes called 'ex ante' pareto superiority – where the 'ex ante' indicates that what are being compared are the expected outcomes for the agents, as calculated *before* the action is carried out. The ex ante pareto principle states that you should not perform an action A_2 if another action A_1 is available such that A_1 is ex ante pareto superior to A_2 . This ex ante pareto principle seems compelling – at least at first blush – and many theorists working in economics and law accept it.⁴ There are, however, theorists who reject it, including Matthew Adler⁵ and Marc Fleurbaey and Alex Voorhoeve (2013).⁶ My argument in this paper undercuts this debate by showing that the very idea of ex ante pareto superiority is (as yet) ill-defined. This is because whether one action is judged to be ex ante pareto superior to another can depend on how the people concerned are designated. Without guidance as to how to designate the people concerned, applications of the concept of ex ante pareto superiority can lead to contradiction. Thus before the concept of ex ante pareto superiority can be coherently applied, its definition must be completed. I explore a variety of ways of completing the definition, and I show that under the most plausible completed version the conditions for application of the concept are more stringent than is usually supposed: some cases of apparent ex ante pareto superiority turn out not to meet the criteria after all. Thus the ex ante pareto principle – with

⁴ For some examples, see footnote 1 of Matthew D. Adler, “The Puzzle of ‘Ex Ante Efficiency’: Does Rational Approvability Have Moral Weight?” *University of Pennsylvania Law Review*, CLI, 3 (January 2003): 1255.

⁵ Adler, *Well-Being and Fair Distribution*, *op. cit.*

⁶ Marc Fleurbaey and Alex Voorhoeve, “Decide as you would with full information!: an argument against ex ante Pareto,” in Nir Eyal, Samia Hurst, Ole Norheim, and Dan Wikler, eds., *Inequalities in Health: Concepts, Measures, and Ethics* (Oxford: OUP, 2013), pp. 113–128.

the completed definition in place – has a more modest reach than is generally supposed, and theorists who have rejected the principle may need to review their stance.

I turn now to argue for my key point: that whether one action A_1 is judged to be ex ante pareto superior to another action A_2 can depend on how the people concerned are designated.

I. WHY IT MATTERS HOW THE PEOPLE CONCERNED ARE DESIGNATED

In assessing whether one action A_1 is ex ante pareto superior to another action A_2 , we consider the *prospects* - i.e. the expected outcome - for the different people involved under each action. The prospects or expected outcome for a person under an action is a function of the outcome for the person under that action at each state, together with the probability of each state. I take the *probability* of the state here to be the epistemic probability from the perspective of the

decision maker.⁷ My key point is that the prospects for an agent can depend on how that agent is designated. Here is a simple example to illustrate this point.

The Jam Competition. Suppose that I am at a fair with my cousin Delia. She has entered the ‘best jam’ competition, and the judges have already tasted the jams and decided on a winner, but they have not yet announced the result. I understand that unfortunately the fair organizers have not been able to afford a prize for the winner of the competition. To pass the time, I try my hand at the tombola, and win a box of chocolates. I don't happen to like chocolate, but it seems that everyone else does. I consider two choices of action: I could give the box of chocolates to my cousin Delia, or I could donate it as the prize for the winner of the best jam competition. I mull over the options:

⁷ A further question is whether we should take the measure of well-being of a person under an outcome to be as assessed by the decision maker: an alternative would be to assess this from the perspective of the person concerned. I don't attempt to settle this question here, but note that however we settle the question of how to measure the well-being of a person under an outcome, we still need some measure of how likely each outcome is under an action, if we are to arrive at the *expected* outcome for each person under an action. My assumption throughout this paper is that the relevant measure of how likely each outcome is under each action is given by the decision maker's epistemic probability function.

A possible alternative is to use the epistemic probability function of each person concerned to calculate his or her own expected outcome under each action. I am sympathetic to this view (despite its numerous difficulties - for which see Adler, “The Puzzle of “Ex Ante Efficiency,” *op. cit.*), and it may be that the key lesson of this paper is that in judging one action to be ex ante pareto superior to another, we must focus exclusively on the prospects as calculated by the people concerned (no matter how ill-informed about the likelihood of the different possible outcomes given various actions) rather than the prospects as calculated using the epistemic probability function of the decision maker.

<i>Action</i>	<i>Prospects (i.e. expected outcome) for my cousin Delia</i>	<i>Prospects (i.e. expected outcome) for the best jam competition winner</i>
GIVE: Give the box of chocolates to my cousin Delia	0.8	0.2
DONATE: Donate the box of chocolates as a prize for the winner of the best jam competition	0.2	0.8

Table 5

If I give the chocolates to my cousin Delia, then I can be pretty sure that Delia will accept them and enjoy eating them – so the prospects for her under this action are good. If I denote the chocolates as a prize for the best jam competition, then the prospects are much less good for Delia: she might be the winner, in which case she'll get the chocolates, but (let's suppose) there are plenty of impressive-looking entries and I don't have much hope that Delia will win. Thus the prospects for Delia are better under GIVE than under DONATE. In contrast, the prospects for the winner of the best jam competition are better under DONATE than under GIVE: the winner of the best jam competition will definitely get the chocolates under DONATE, but only get them under GIVE if Delia turns out to be the winner.

Of course the twist in this tale is that my cousin *is* the winner of the best jam competition, though neither she nor I know it yet. Thus whichever action I take, the outcome will be the same for my cousin Delia as it is for the winner of the best jam competition. Nevertheless, for each action the *prospects* (i.e. the *expected* outcome) for Delia and the winner of the best jam competition are different. This is because the prospects for the people concerned under my possible actions depend on my (the decision maker's) epistemic state, and I do not know that

the winner of the best jam competition is my cousin Delia.⁸ This illustrates the fact that the prospects for a person under an action can depend on how that person is designated. Another way to put this point is to say that here we have an intensional or opaque context: the truth value of the claim that a particular action has particular prospects for a person can alter when the very same person is designated differently.⁹

Having argued that the prospects under an action for an agent can depend on how that agent is designated, I now draw out the implications of this point for the concept of ex ante pareto superiority.

II. THE CONCEPT OF EX ANTE PARETO SUPERIORITY IS ILL-DEFINED

An action A_1 counts as pareto superior to another action A_2 iff for each person concerned, the prospects for that person under action A_1 are at least as good as the prospects for that person under action A_2 , and furthermore for at least one person, the prospects for that person under action A_1 are better than the prospects for that person under action A_2 . We have seen that the

⁸ This line of argument implies (correctly, on my view) that “referentialism about credence” (David J. Chalmers, “Frege’s Puzzle and the Objects of Credence,” *Mind*, CXX, 479 (November 2011): 587–635, at p. 590) – or at least “naive referentialism about credence” (Jesse Fitts, “Chalmers on the Objects of Credence,” *Philosophical Studies*, CLXX, 2 (October 2013): 343–358) – is false.

⁹ This holds even if we limit our attention to rigid designators. To see this, we can replace ‘the winner of the best jam competition’ throughout with ‘the *actual* winner of the best jam competition’ (thereby rigidifying the description); or we can suppose that the judges have nicknamed the winner ‘Mrs Beaton’ (because the winning jam reminds them of a recipe from Mrs Beaton’s book), and I have overheard them unanimously agreeing that Mrs Beaton is the winner – in which case we can replace ‘the winner of the best jam competition’ throughout with ‘Mrs Beaton’.

prospects for a person under an action can depend on how that person is designated. Thus the definition of pareto superiority is incomplete: how should we designate the people concerned when calculating their prospects? If we ignore this question and assume that the definition works just fine regardless of how we designate the people concerned, then we get into trouble, for applications of the concept of pareto superiority can give us inconsistent results.

Here is an example to illustrate this. It is a stylized and unrealistic example, but it is useful for seeing the point clearly.¹⁰

Disease X and disease Y. Suppose that you are a doctor, and you are told that there are two patients in a single ward, both in need of urgent medical help. One has deadly disease X, and the other has deadly disease Y. You have two canisters of gas – one designed to treat disease X (which nearly always completely cures disease X, almost certainly has no effect on disease Y, and causes no side-effects) and one designed to treat disease Y (which nearly always completely cures disease Y, almost certainly has no effect on disease X, and inevitably causes permanent blindness as a side-effect). There is no time to move either patient out of the ward, and only time to pump in one canister of gas. Thus you must either pump in the gas designed to cure disease X (X-CURE), thereby (almost certainly) curing one patient while leaving the other to die, or you must pump in the gas designed to cure disease Y (Y-CURE), thereby leaving one patient (almost certainly) alive but permanently blind, and the other dead. The table below gives the prospects for each person concerned (i.e. each patient) under each action:¹¹

¹⁰ This scenario is inspired by examples from Caspar Hare, “Take the Sugar,” *Analysis*, LXX, 2 (January 2010): 237–247, and Caspar Hare, *The Limits of Kindness* (Oxford: OUP, 2013).

¹¹ Note that – because it is only *almost* certain that each cure will work on the disease it is designed to treat – you are here comparing the *prospects* for the two patients, rather than the outcomes.

	<i>Prospects for the patient with disease X</i>	<i>Prospects for the patient with disease Y</i>
X-CURE	0.9	0
Y-CURE	0	0.7

Table 6

We can see at once that neither action is ex ante pareto superior to the other. X-CURE is better for the patient with disease X, and Y-CURE is better for the patient with disease Y. You might decide that one action is better all things considered, but it cannot be on the basis that it is the ex ante pareto superior action.

Now suppose that you rush down the corridor to the ward with your canisters of gas, wondering which you ought to pump in when you arrive. When you reach the ward, you see two patients lying in bed – both unconscious. One has notes identifying him as Ethan, and the other has notes identifying him as Fred. They look like identical twins. You have no way of telling quickly who has disease X and who has disease Y, because these diseases have similar symptoms. Your epistemic probability that Ethan has disease X and Fred has disease Y is 0.5, and similarly your epistemic probability that Ethan has disease Y and Fred has disease X is 0.5. Now you calculate again the prospects of your actions for the two patients:

	Ethan	Fred
X-CURE	$(0.5*0.9) + (0.5*0) = 0.45$	$(0.5*0.9) + (0.5*0) = 0.45$
Y-CURE	$(0.5*0.7) + (0.5*0) = 0.35$	$(0.5*0.7) + (0.5*0) = 0.35$

Table 7

Action X-CURE emerges as ex ante pareto superior. The prospects for both Ethan and Fred are better under action X-CURE than under action Y-CURE.

We can see then that whether one action is classed as ex ante pareto superior to another can depend on how the people concerned are designated. If you designate the two patients as Fred and Ethan, and consider *their* prospects, then action X-CURE comes out as ex ante pareto superior to action Y-CURE. However if you designate the same two patients instead as the patient with disease X and the patient with disease Y, and consider *their* prospects, then action X-CURE does not come out as ex ante pareto superior to action Y-CURE. Thus the concept of ex ante pareto superiority is ill-defined as it stands. If we attempt to apply it without completing its definition, then it can generate inconsistent results.¹²

III. COMPLETING THE CONCEPT OF EX ANTE PARETO SUPERIORITY

How then should we complete the definition of pareto superiority? Let us start by re-stating the incomplete criterion for action A_1 to count as pareto superior to action A_2 : for every person concerned, the prospects under A_1 are at least as good as the prospects under A_2 , and for at

¹² It may be objected that there is a natural and obvious way to apply the concept of ex ante pareto superiority in any relevant scenario, for in any relevant scenario there will be some natural and obvious way to designate each of the people concerned. My first response is to point out that there are scenarios in which there is no *one* natural and obvious way to designate the people concerned, and my example of *Disease X and Disease Y* is an instance of this. However it might be countered that this scenario is gerrymandered, and real-life scenarios with similar features are rare: thus though we do still need to complete the concept of ex ante pareto superiority if we want it to apply quite generally, this is not an urgent task given that in most real-life scenarios where we wish to apply the concept we can do so using the natural and obvious way to designate the people concerned. In response, I would argue that in *any* scenario there is a range of possible ways to designate the people concerned, and so a choice to be made over how the concept of ex ante pareto superiority is to be applied. One choice may seem natural or obvious, but it is worth examining the reasons for that choice – for the choice may have serious consequences: whether one action is judged ex ante pareto superior to another may depend on this choice of designators. Thus I maintain that the concept of ex ante pareto superiority is incomplete as it stands, and that it does not apply to any scenario until a choice has been made about how it ought to be completed. Many thanks to an anonymous reviewer for pushing me on this question.

least one person, the prospects under A_1 are better than the prospects under A_2 . There are (at least) three ways in which we might try to extend this into a complete definition:

1. Specify a particular way of designating people: iff the criterion is met with the people concerned designated in *this* way, then action A_1 is pareto superior to action A_2 .
2. Require only that the criterion must be met for *some* way of designating all the people concerned: iff this holds, action A_1 is pareto superior to action A_2 .
3. Require that the criterion must be met for *all* ways of designating all the people concerned: iff this holds, action A_1 is pareto superior to action A_2 .

I shall now explore each of these options.

1. The criterion must be met for some special way of designating the people concerned.

The challenge here is to find a special way of designating people that can do the job. It must be unique – or at least, if a person can have two designators, both of which designate him or her in the special way – then it must be impossible for a rational decision maker not to know that these two designators co-refer. Otherwise the concept of ex ante pareto superiority will still be ill-defined.

To illustrate the difficulty, here is one attempt to carve out a category of special designators. We might draw a distinction between designators which are opaque and designators which are

transparent.¹³ The test of whether a designator is opaque or transparent is this: does the decision maker know to whom the designator refers? If so, then the designator is transparent; if not, then the designator is opaque. Perhaps in the case of *Disease X and disease Y*, the designators 'Ethan' and 'Fred' are transparent, while the designators 'the patient with disease X' and 'the patient with disease Y' are not – because after all you don't know to whom 'the patient with disease X/Y' refers. The trouble here is that the distinction between opaque and transparent designators is elusive, and when we attempt to spell it out, it becomes clear that it cannot do the job required. To see this, consider first that in a sense of course you do know to whom 'the patient with disease X' refers: it refers to *the patient with disease X*. What else is needed to make this designator transparent? Do you need to know some proper name of the referent? Or some identifying fact about the person's time-space location (e.g. that the person is/was at a particular location at a particular time)? Or what? For any of these requirements, we can adjust the case of *Disease X and disease Y* so that the designators 'the patient with disease X/Y' are transparent after all. For example, we can suppose that the two men have been your patients for a long time, and you have met them individually plenty of times before. You have never bothered to learn their real names, but you and the other hospital staff refer privately to one as 'Grumpy' and the other as 'Whiney'. When you arrive at the hospital, you learn that Grumpy has disease X, and Whiney has disease Y. Of course then you do know to whom 'the person with disease X' refers (it refers to one of your oldest patients – Grumpy!), and to whom 'the person with disease Y' refers (it refers to good old Whiney!): these are then transparent rather than opaque designators. This can hold even if when you reach the ward, the patients are so ill and immobile that you cannot tell by sight which is Grumpy and which is Whiney – but rely on the notes to

¹³ Hare uses this terminology – though for cases rather than designators. Hare counts a case as 'fully transparent' only if the decision maker knows everything about the people concerned – but he also discusses 'semi-transparent' cases, and this is a category to which some of the cases I discuss belong (Hare, *The Limits of Kindness*, *op. cit.*, p. 181).

identify the two patients as Ethan and Fred. Thus there are two transparent designators for each agent, and the prospects for each agent depend on which set of transparent designators are employed. The requirement that the designators be transparent then is not enough to complete the definition of pareto superiority.

An alternative suggestion is that the special way of designating the people concerned is by proper name.¹⁴ Thus in the original case of *Disease X and disease Y*, it is the prospects for Fred and Ethan that are relevant (because 'Fred' and 'Ethan' are proper names). The prospects for both patients (under these designators) are better under X-CURE than under Y-CURE, and so (on this view) X-CURE is ex ante pareto superior to Y-CURE. Of course, the prospects for the patient with disease Y are worse under X-CURE than under Y-CURE, but this is irrelevant because 'the patient with disease Y' is not a proper name. The problem with this suggestion is that individuals can have more than one proper name, and even a rational decision maker may not know that the names co-refer. Thus in the adjusted version of the scenario *Disease X and disease Y* described above, there are two sets of proper names for the patients: 'Ethan' and 'Fred', and 'Grumpy' and 'Whiney'.¹⁵ Though X-CURE meets the criteria for pareto superiority over Y-CURE when the prospects for the two patients are as calculated under their designators 'Ethan' and 'Fred', X-CURE fails the criteria when the prospects for the two patients are calculated under their designators 'Grumpy' and 'Whiney'. Thus whether X-CURE is ex ante pareto superior to Y-CURE depends on how the people concerned are designated: in this case, by which *proper names* they are designated. Thus specifying that proper names are the relevant

¹⁴ Thanks to an anonymous reviewer for pushing this point.

¹⁵ I class 'Grumpy' and 'Whiney' as proper names - albeit nicknames - but note that the argument would go through in the same way if we substituted 'Mr Smith' and 'Mr Jones' for these names throughout.

special designators does not solve the problem: the concept of ex ante pareto superiority remains ill defined.

Are there other special designators that could do the job? We might try focusing on a person's full legal name, rather than on all his or her proper names (which may include nicknames). Usually a person has no more than one full legal name,¹⁶ so perhaps by focusing on full legal names we can avoid the problems that we faced above when we focused on proper names more generally. But the problem now is that not everybody has a full legal name: for example, in many countries babies are not required to be legally named until they reach a certain age, and so many newborn babies lack full legal names. And yet of course whether one action is ex ante pareto superior to another may depend on the prospects for a baby – regardless of whether it has a full legal name or not. Thus we cannot complete the concept of ex ante pareto superiority by requiring that the people concerned be designated by their full legal names, and expect the concept to apply quite generally. The same goes for other civil identifiers, such as passport numbers or national insurance numbers: there will be people whose prospects matter who cannot be designated in this way.¹⁷

To make this approach work, then, we need to carve out some category of 'special' designators that can do the job – and I do not see how this can be done.

2. The criterion must be met for some way of designating the people involved. With this option, to demonstrate that one action is ex ante pareto superior to another, you need only find *some way* of designating each of the people involved, such that so designated the criterion is

¹⁶ Although this does not always hold when a person has dual citizenship.

¹⁷ Many thanks to an anonymous reviewer for suggesting the idea of a 'civil identity'.

met. This would mean that in the case of *Disease X and disease Y*, X-CURE is the pareto superior action. This is because if we designate the two people concerned as Ethan and Fred, then each of them (so designated) have better prospects under X-CURE than under Y-CURE. Thus the criterion for pareto superiority is met for *some* way of designating all the people involved, and – on the option that we are exploring – this is all that is required for pareto superiority.

The problem with this option is that it becomes all too easy to show that one action is pareto superior to another. Effectively, whenever one action produces a greater amount of well-being overall relative to some other action, the first action will be pareto superior to the second – regardless of how the well-being is distributed. To see this, consider again the scenario at the start of this paper, *Two Rabies Patients*:

	<i>Outcome for Alan</i>	<i>Outcome for Bob</i>
Give the medicine to Alan	0.9	0
Give the medicine to Bob	0	0.85

Table 1

Here the rabies medicine would save either Alan or Bob from certain death (let us suppose), but it would produce slightly worse side-effects for Bob. This seems like a clear case where neither action is pareto superior to the other: giving the medicine to Alan is better for Alan, but giving the medicine to Bob is better for Bob. However, there are other designators available for the people involved. For example, one patient presumably cut his or her first tooth at a younger age than the other: let us call the person who cut his tooth first the 'first toother', and the other person the 'second toother'. You have no idea whether Alan is the first toother and

Bob the second toother, or vice versa: you have an epistemic probability of 0.5 in each of these two possibilities.

If you give the medicine to Alan, say, what are the prospects for the first toother? Well, you have a 0.5 epistemic probability that the first toother is Alan, in which case he'll get the medicine and it will cure him (and his level of well-being will be 0.9); and you have a 0.5 epistemic probability that the first toother is Bob, in which case he won't get the medicine and will die (in which case his level of well-being will be 0). From this we can calculate the prospects for the first-toother under the action of giving the medicine to Alan, as $0.5 \cdot 0.9 + 0.5 \cdot 0 = 0.45$. The prospects for the first and second tooters under your actions are set out in the table below:

	Prospect for First-Toother	Prospect for Second-Toother
Give the medicine to Alan	$0.5 \cdot 0.9 + 0.5 \cdot 0 = 0.45$	$0.5 \cdot 0.9 + 0.5 \cdot 0 = 0.45$
Give the medicine to Bob	$0.5 \cdot 0.85 + 0.5 \cdot 0 = 0.425$	$0.5 \cdot 0.85 + 0.5 \cdot 0 = 0.425$

Table 8

We can see then that giving the medicine to Alan (as opposed to giving the medicine to Bob) has better prospects for the first toother and better prospects for the second toother – and these are the only two people concerned. Thus the criterion for ex ante pareto superiority is met under *some* way of designating the people concerned (namely as 'first toother' and 'second toother') and so, under the option for completing the definition of pareto superiority that we are exploring, giving the medicine to Alan is ex ante pareto superior to giving the medicine to Bob.

In general, whenever the total expected well-being produced by one action is greater than the total expected well-being produced by another action, the first will be ex ante pareto superior

to the second. To see why this is, consider that we can typically find an uninformative and whimsical way of designating the people concerned, so that the decision maker cannot tell how each of the people (so designated) will be affected by his or her actions: the decision maker only knows that the people so designated match up somehow with the people concerned. So designated, the prospects for each person under a given action will be the same: we simply take the total expected outcome of the action, and divide it up equally to get the prospects for each person (so designated). Thus whenever an action A_1 has a higher expected outcome overall (regardless of how it is distributed) than an action A_2 , each person – under their whimsical and uninformative designator – will have better prospects under action A_1 than under action A_2 . Thus a higher overall expected outcome will automatically guarantee ex ante pareto superiority.

Under this option, ex ante pareto superiority collapses into a coarser measure of superiority: the *distribution* of prospects do not play a role – it is just the overall expected prospects that matter. This is entirely opposed to the spirit of the concept, so I conclude that this option for *completing* the definition of ex ante pareto superiority is unsuccessful, and turn to our final option.

3. *The criterion must be met for every way of designating the people involved.* This option strikes me as the most plausible way to complete the definition of ex ante pareto superiority. Nevertheless, it faces some difficulties. Thus the reader is left with a choice: the whole concept of ex ante pareto superiority can be dropped as incoherent; or this way of completing the definition can be developed and defended. I begin by explaining the concept with the definition completed in this way, fleshing it out to make it as appealing as I can. Then I turn to consider some of the implications and difficulties of accepting it.

I start by clarifying what is meant by *every* way of designating the people involved: what is the relevant domain here? In order to ensure that the concept of ex ante pareto superiority can play a role in guiding decisions, I will take the relevant domain of designators to be those that the decision maker recognizes as designating the people concerned. Thus for example in the original *Disease X and disease Y* scenario, the domain will include at least 'Ethan', 'Fred', 'the patient with disease X', 'the patient with disease Y', and no doubt many other designators. But if the patient with disease X happens to be secretly terminally ill and suicidal, then the designator 'the terminally ill and suicidal patient at this hospital with disease X' will not be included in the domain, because you (the decision maker) are not sure that this designates anybody. For an action A_1 to count as ex ante pareto superior to an action A_2 , then, the prospects under A_1 must be at least as good as the prospects under A_2 for every person under every designator in the domain; and the prospects under A_1 must be better than the prospects under A_2 for at least one person under some designator in the domain.

Under this option for completing the definition of ex ante pareto superiority, the conditions for ex ante pareto superiority are more stringent than is usually assumed. It might seem that they are so stringent they can never be met – or at least that they can never be shown to be met. After all, to establish that one action is ex ante pareto superior to another, the decision maker has to consider every designator that (s)he recognizes as designating one of the people concerned, and calculate the prospects under each action for each person so designated. How could this be accomplished in a finite length of time? I agree that a *proof* of ex ante pareto superiority is hard to produce under this option, but that need not mean that ex ante pareto superiority cannot be a useful concept in decision making. Here is one possible decision

programme for cases where you are faced with the choice between two actions, A_1 and A_2 .¹⁸ First consider whether there is anyone – under any designator – whose prospects are *worse* under action A_1 than they are under action A_2 . If you can think of some such designator, then you have established that action A_1 is not ex ante pareto superior to action A_2 . Whether you go on to perform action A_1 will depend on your decision theory more generally – but at any rate you cannot perform action A_1 on the grounds that it is ex ante pareto superior to action A_2 . If, on the other hand, you cannot think of any such designator, then you may proceed on the assumption that this necessary condition for ex ante pareto superiority is met. Now it merely remains to show that action A_1 has better prospects for at least one person under one designator, and once this is accomplished you may take it for granted that action A_1 is pareto superior to action A_2 .

To illustrate this strategy, consider first the scenario *Disease X and disease Y*. Releasing the cure for disease X (rather than the cure for disease Y) has better prospects for both Ethan and Fred – and 'Ethan' and 'Fred' designate the only two people concerned by your action. However, this does not establish that releasing the cure for disease X (rather than disease Y) is ex ante pareto superior: you must cast around for any designator under which someone stands to lose from your choice of action. And 'the person who has disease Y' is such a designator: the prospects for the person who has disease Y are worse if you release the disease X cure, rather than the disease Y cure. This establishes that releasing the cure for disease X is not ex ante pareto superior to releasing the cure for disease Y: you might still choose to release the cure

¹⁸ It is easy to see how this decision programme could be extended to cases where your choice of actions is larger.

for disease X, of course, but you cannot do so on the grounds that it is ex ante pareto superior to all the other available actions.

Let us now compare the scenario *One Rabies, One Healthy*.¹⁹ In this scenario, giving the rabies medicine to Alan (rather than to Carol) has both better prospects for Alan and better prospects for Carol. Is there any way of designating any of the people concerned so that his or her prospects are worse if the medicine is given to Alan rather than Carol? Unless you can think of any (and I can't) you can safely assume that giving the medicine to Alan is pareto superior, and (depending on your decision theory more generally) you may give the medicine to Alan on the grounds that it is ex ante pareto superior to any other action available.

Having explained and illustrated the concept of ex ante pareto superiority completed in this way, I now draw out some connections with other work in the literature, and consider some implications and difficulties for the concept so completed.

IV. IMPLICATIONS FOR A RANGE OF DEBATES

The ex ante pareto principle plays an important role in various debates. Not everyone has accepted the principle, but it certainly seems to have an intuitive pull, and those who reject it are assumed to have some explaining to do. We now need to reappraise this dialectical situation in the light of the following facts: the concept of ex ante pareto superiority is ill-defined as it stands, and under the most plausible way of completing the definition, the conditions for ex ante pareto superiority are more stringent than generally thought. These two points have

¹⁹ This works either with the simplest version (in which you can be certain of the outcomes) given in table 2, or the more complex version given in table 3.

significant repercussions for various debates. As an illustration, I discuss the impact on Adler's argument for his social welfare function.²⁰

Adler argues for a particular way of ranking outcomes, and his ranking is prioritarian – meaning (very roughly) that the well-being of the worst off under some outcome is a weighty consideration in assessing that outcome as a whole. Adler then merges this ranking of outcomes with a version of Expected Utility Theory, to give a ranking of actions (or choices, in Adler's terminology). His ranking of actions leads to apparent violations of the ex ante pareto principle. To see this, consider the following case (a rough simplification of the case discussed by Adler:²¹

	<i>State S₁</i> $P(S_1) = 0.5$		<i>State S₂</i> $P(S_2) = 0.5$	
	<i>Outcome for Gregor</i>	<i>Outcome for Holly</i>	<i>Outcome for Gregor</i>	<i>Outcome for Holly</i>
Action A ₁	0.9	0.1	0.1	0.9
Action A ₂	0.5-ε	0.5-ε	0.5-ε	0.5-ε

Table 9

The expected outcome for Gregor under action A₁ is $(0.5*0.9)+(0.5*0.1) = 0.5$, and similarly, the expected outcome for Holly under action A₁ is 0.5, whereas the expected outcome for each of these people under action A₂ is 0.5-ε. Thus both Gregor and Holly have better prospects

²⁰ Adler, *Well-being and Fair Distribution*, *op. cit.*

²¹ *Ibid.* p. 504.

under action A_1 than they have under action A_2 , and so it seems that action A_1 is ex ante pareto superior. Thus, by the ex ante pareto principle, action A_1 is the morally superior action.

However by Adler's ranking, provided that ϵ is sufficiently small, action A_2 is morally superior to action A_1 . For the outcome of action A_2 is ranked higher than either of the possible outcomes of action A_1 : the gain in well-being for the worst off (from 0.1 to $0.5-\epsilon$) outweighs the loss in well-being for the better off (from $0.9 - 0.5-\epsilon$) – despite the fact that the worst off person's gain is smaller than the better off person's loss. Thus Adler's view appears to conflict with the ex ante pareto principle, and Adler acknowledges that this is a 'troubling feature' of his view.²² In response, Adler demonstrates that alternative theories in the same vein display even more troubling features.

But if the concept of ex ante pareto superiority is ill-defined, then the ex ante pareto principle is not true (or even meaningful). There should then be no pressure to accommodate the principle, nor any need to justify the rejection of it. We first need to complete the definition of ex ante pareto superiority and then consider the principle that would result. If we complete the definition in the way that I recommend in the previous section, then it is no longer obvious that the principle conflicts with Adler's account. To see this, consider again the table above, with the choice between action A_1 and action A_2 . It is true that the prospects for both Gregor and Holly are better under action A_1 than under action A_2 , but to establish that action A_1 is pareto superior to action A_2 , we need also to consider other ways of designating the individuals involved. What about 'the person who will lose out if action A_1 is performed, rather than action A_2 '? Whether this designates an individual might depend on the details of the scenario (and see the discussion below under section V), but one way to fill in the scenario is as follows. The

²² Ibid. p. 504.

names 'Gregor' and 'Holly' have been written on slips of paper, and one has been drawn at random – though the decision maker does not know yet which name has been drawn. In state S_1 , Holly's name has been drawn, and on that basis she gets the bad outcome; in state S_2 , Gregor's name has been drawn, and on that basis he gets the bad outcome. What are the prospects for the person whose name has been drawn?²³ If the decision maker performs action A_1 , then the prospects (i.e. the expected outcome – as calculated by the decision maker) for this person will be 0.1. If the decision maker performs action A_2 , then the prospects for this person will be $0.5-\epsilon$. Thus it is not the case that each person *under every designator* has prospects that are at least as good under action A_1 as they are under action A_2 , and so action A_1 is not ex ante pareto superior (on my recommended way of completing this concept) to action A_2 . Thus Adler's claim that action A_2 is the morally superior action does not conflict with the ex ante pareto principle here.

As we shall see, there may be ways of completing the scenario so that action A_1 does emerge as ex ante pareto superior: much depends on the details of the scenario, and niceties concerning the notion of a designator. At any rate I take the dialectical situation to be weighted much more in Adler's favour than he thinks. The concept of ex ante pareto superiority is ill-defined as it stands – and so the ex ante pareto principle is similarly ill-defined and can hardly be compelling. When we attempt to complete the concept in the most plausible way, we end up with a principle which has not yet been shown to conflict with Adler's social welfare function. Whether Adler has an objection to answer to here is not yet settled.

²³ To allay any worries that 'the person whose name has been drawn' is a non-rigid designator, we can substitute this description with 'the person whose name has *actually* been drawn' – or coin a name for the person ('Unlucky').

Adler is not the only theorist who rejects the ex ante pareto principle: I have just taken his theory as an illustration.²⁴ The general point here is that once it is seen that the definition of pareto superiority is incomplete as it stands, and how it might be developed into a complete definition, there are implications for a range of debates.

V. A PROBLEM FOR THE EX ANTE PARETO PRINCIPLE

In this section I discuss a problem for the ex ante pareto principle, with the definition of ‘pareto superiority’ completed as I recommend. This problem did not originate with the completion of the definition and has already been discussed in the literature. Completing the definition of pareto superiority solves one problem for the ex ante pareto principle (namely, it renders it well-defined), but it does not solve the problem I discuss in this section. The problem is that the moral goodness of an action does not seem to depend just on the prospects for *individuals*.

²⁴ Another interesting and recent example can be found in Fleurbaey and Voorhoeve, “Decide as You Would with Full Information!”, op.cit. Fleurbaey and Voorhoeve argue that if an egalitarian principle is combined with their plausible “Principle[s] of Full Information”, violations of the ex ante pareto principle result. They use this point to argue that the ex ante pareto principle is false. In contrast, I have argued in this paper that the concept of ex ante pareto superiority (as it is generally discussed in the literature, and as it is expressed by Fleurbaey and Voorhoeve) is ill-defined, and so that the resulting ex ante pareto principle is not even meaningful, and so cannot conflict with anything or be shown to be false. I have argued that on the only plausible way of completing the concept, we get an ex ante pareto principle with a more modest reach than is generally assumed. The resulting principle does not conflict with Fleurbaey and Voorhoeve's Principles of Full Information combined with egalitarian principles – at least not in the examples that they discuss. This is a further reason to accept the ex ante pareto principle under my recommended completion of the definition: the resulting principle is compatible with egalitarianism together with the plausible Principles of Full Information. Thanks to an anonymous reviewer for drawing my attention to these connections.

And this holds even if we understand the prospects for the individuals concerned to include their prospects under all the relevant designators. It seems that we are also concerned with the prospects for 'merely statistical people' (Hare, 2012). To see the issue here, I turn again to my rough version of Adler's scenario, but elaborate the case differently.

Recall that in this scenario, if the decision maker carries out action A_1 , then either the outcome for Gregor will be 0.9, and that for Holly 0.1, or vice versa (with the decision maker's probability equally divided between these possibilities); if instead the decision maker carries out action A_2 , then the outcome for both Gregor and Holly will be a sure $0.5-\epsilon$. Earlier I elaborated this scenario by stating that the name of the person who would get the bad outcome under action A_1 had already been randomly selected before the decision maker chooses between A_1 and A_2 . Thus 'the person whose name has been drawn' denotes one of the individuals concerned. The prospects for this person so designated are worse under action A_1 than they are under action A_2 ; this shows (under my preferred way of completing the concept) that action A_1 is not ex ante pareto superior to action A_2 . But now let us elaborate the scenario differently. Let us suppose that if the decision maker chooses to carry out action A_2 , then no name will be selected; if the decision maker chooses to carry out action A_1 , then and only then will a name get randomly pulled from the hat to decide who will get the bad outcome. In this case, can we find any way of designating an individual whose prospects are worse under action A_1 than they are under action A_2 ? We might try 'the person whose name would be pulled from the hat should the decision maker choose action A_1 ' – but the worry is that this may not designate an individual at all. If the decision maker in fact chooses action A_1 , then a name will be pulled from the hat – and so plausibly the designator picks out the person so named. But if instead the decision maker chooses action A_2 , then is there a person whose name *would* have been pulled from the hat had action A_1 been chosen? The closest worlds at which action A_1 is

chosen will include both worlds where Gregor's name is pulled from the hat, and worlds where Holly's name is pulled from the hat – and so it seems that designator does not uniquely designate either person. We might say that it designates an 'indeterminate' person – akin to a “fictional character”²⁵ – or a “statistical person”.²⁶

There is evidence to suggest that we are on the whole more motivated by concern for particular individuals than by concern for merely statistical people: this is the ‘identifiable victim effect’, and it is invoked to explain why it was possible to raise hundreds of thousands of dollars in 1987 to rescue Jessica McClure from a well, but much more difficult to raise money for preventative healthcare which would save the lives of hundreds of unidentified children.²⁷ But even if we indeed ought to prioritise the well-being of the identified over the merely statistical (which of course does not necessarily follow from the fact that we do seem to prioritise in this way), concern for the well-being of the merely statistical person may still carry *some* weight. Hare, for example, argues persuasively that while our concern for the well-being of an identified victim ought perhaps to carry more weight than our concern for the well-being of a merely statistical victim, it should not carry twice as much weight.²⁸

If the prospects for merely statistical people ought to carry some weight, then perhaps we should extend our definition of pareto superiority to include merely statistical people as well

²⁵ Michael Otsuka, “Risking Life and Limb: How to Discount Harms by Their Improbability,” in Glenn Cohen, Norman Daniels, and Nir Eyal, eds., *Identified versus Statistical Lives: An Interdisciplinary Perspective* (Oxford: OUP 2015), p. 85.

²⁶ Caspar Hare, “Obligations to Merely Statistical People,” *The Journal of Philosophy*, CIX, 5/6 (May/June 2012): 378-390.

²⁷ Karen Jenni and George Loewenstein, “Explaining the Identifiable Victim Effect,” *Journal of Risk and Uncertainty*, XIV, 3 (May/June 1997): 235-257.

²⁸ Hare, “Obligations to Merely Statistical People”, *op. cit.*

as individuals. This would require a good deal of care. If our concern is not limited to prospects for *people* – then what limits should be in place? It may be that the challenge of reforming the concept of ex ante pareto superiority to encompass prospects for statistical as well as actual people is too great, and the concept ought to be abandoned all together as unfit for purpose. This paper does not attempt to solve this problem: indeed, it brings it into sharper focus.

VI. CONCLUSION

I have argued that the concept of ex ante pareto superiority is ill-defined as it stands. For whether one action is ex ante pareto superior to another depends on the prospects for each person concerned under each action, and these prospects in turn depend on how the people concerned are designated. Without any restrictions on how the people concerned should be designated, applications of the concept give inconsistent results. I have explored various options for completing the definition (and so rendering the concept well-defined), and I have argued that on the most plausible option, one action is classed as pareto superior to another only if the condition is met under *every* relevant way of designating the people concerned. The result is that the relation of ex ante pareto superiority holds in fewer cases than previously thought, and consequently the ex ante pareto principle has a surprisingly modest reach.