

Decide as You Would with Full Information! An Argument against *ex ante* Pareto¹

Marc Fleurbaey and Alex Voorhoeve

As policy-makers and private individuals, we sometimes face choices of the following kind: either choose an alternative which has somewhat lower *expected* value for each person, but which will substantially improve the *outcomes* of the worst off, or choose an alternative which has somewhat higher expected value for each person, but which will leave those who end up worst off substantially less well off. By way of illustration, consider the following real-world MAMMOGRAM CASE (U.S. Preventative Services Task Force, 2009):

In 2009, the U.S. Preventive Services Task Force had to decide whether or not to recommend routine biannual breast cancer screening mammography for all women aged 40 to 49. For each woman in this age group, the potential costs of this method of screening include the small to moderate psychological and physical harms caused by initial testing and by follow-up testing.² Its potential benefits are improvements in

¹ Earlier versions were presented at the Brocher Foundation, Carnegie Mellon University, Groningen University, the Hebrew University, Leeds University, the LSE, the NIH, Oxford University, UCL, and the University of Pennsylvania. We thank audience members for helpful discussions. We also thank Matthew Adler, John Broome, Boudewijn de Bruin, Nir Eyal, Johann Frick, Ipek Gençsü, Alex John London, Jerry Menikoff, Serena Olsaretti, Shepley Orr, Matthew Rendall, Jesse Saloom, Shlomi Segall, Peter Vallentyne, Andrew Williams, and two anonymous referees for comments. We are especially grateful to Michael Otsuka for lengthy discussion and detailed comments.

² For simplicity, we assume the harms do not include an increased risk of breast cancer due to exposure to radiation in testing. This risk is thought to be very low (U.S. Preventative Services Task Force 2009, p. 271).

length and quality of life due to the early diagnosis and more successful treatment of cancer.

The Task Force judged that for each woman in this age group, the expected costs of screening outweighed the expected benefits by a modest margin.³ The alternative we will call ROUTINE SCREENING would therefore have a somewhat lower expected value for each woman in this group, but considerably improve the lot of those among them who, because of undiagnosed cancer, would otherwise be among the worst off. The alternative we will call NO ROUTINE SCREENING would have slightly higher expected value for each woman, but would lead to more early deaths and greater inequality in health outcomes.

In this chapter, we analyse how one should decide such cases in which the alternative that has the highest expected value for each person also disadvantages those who end up worst off. One principle which is commonly invoked in such cases is *ex ante* Pareto. This principle holds that if an alternative has higher expected utility for every person than every other alternative, then this alternative should be chosen. Indeed, the Task Force was moved by this principle: it regarded the expected disvalue that ROUTINE SCREENING would have for each woman as decisive, and recommended against it. Notwithstanding the widespread support of *ex ante* Pareto, we shall argue that, in a certain type of case, it should be rejected because it conflicts with a requirement of rationality. This requirement is that when one lacks information, but can infer that there is a particular alternative one *would* invariably regard as best if one had full information, then one should choose this alternative.

³ This recommendation did not apply to women who are known to be at atypically high risk of developing breast cancer. For simplicity, our discussion excludes this sub-group.

Our argument proceeds as follows. Real-world decision problems like the MAMMOGRAM CASE typically involve a multitude of morally relevant factors. In order to focus our analysis on the key issue, we develop a stylized case in Section I. In Section II, we introduce our rationality principle and show that it is inconsistent with *ex ante* Pareto in our stylized case. In Section III, we extend this rationality principle. In Section IV, we show that it does not rule out other *ex ante* considerations, such as a concern for the fairness of equal-chance lotteries. In Section V, we discuss the relevance of our conclusions for real-world cases, including the MAMMOGRAM CASE.

I. A Stylized Case

The MAMMOGRAM CASE involves two morally relevant considerations that are extraneous to our core question. First, even if the Task Force had decided in favour of routine mammography, this would merely have involved advising women to have mammograms, coupled with steps to ensure these mammograms were accessible to all. The choice of whether or not to have a mammogram would still have been up to each woman. Second, the officials on the Task Force or the doctors who would be acting on its advice may be regarded as agents of the women affected.⁴

In order to focus on the conflict between *ex ante* Pareto and a concern for those who will end up worst off, we shall therefore devote much of this chapter to an analysis of the following stylized case in which there is no choice on the part of the affected individuals and

⁴ A further consideration is that testing women may lead to harm, whereas not testing them merely makes it less likely that one can do good for them. For simplicity, we set aside this aspect of the MAMMOGRAM CASE. Note that our stylized cases do not involve the imposition of harm.

the benefactor is not an agent of the affected individuals, but merely a morally motivated stranger.

VISUAL IMPAIRMENT CASE 1:

Imagine that two ten year-old children, Adam and Bill, have excellent vision but will soon go totally blind due to natural causes unless a morally motivated stranger, Teresa, intervenes. Teresa can use a resource she rightfully controls to produce and administer only one of the following two medicines to both Adam and Bill.⁵

The Egalitarian Medicine (M_E) will ensure that Adam and Bill will both be equally, and considerably, visually impaired, though not totally blind: they will be able to make out basic shapes at a moderate distance, but even with glasses, they will be unable to recognize a friend on the other side of the street and be able to read ordinary newsprint only with difficulty.⁶

The *ex ante* Pareto-optimal Medicine (M_P) will either, with 50 percent probability, cause Adam to retain excellent vision and Bill to have a severe visual impairment (which differs from total blindness only in that one can make out basic shapes at a short distance) or instead cause Bill to retain excellent vision while leaving Adam with the severe visual impairment.

⁵ We assume those in need of aid are children in order to eliminate the consideration that a morally-motivated stranger has reason to do what will maximize the expected utility of each out of respect for their autonomy. (One might argue that a respect for autonomy would require such expected utility-maximization if one assumed a preference-based measure of utility of the kind elaborated below.) For further discussion, see Otsuka and Voorhoeve (2009, 185-8).

⁶ These descriptions of visual impairments and the associated levels of well-being are derived, with some modifications, from the well-known Health Utilities Index-3. See www.healthutilities.com.

We will now make some assumptions about the utilities associated with these treatments. Throughout, we will assume a measure of utility that is derived from idealized preferences satisfying the Von Neumann-Morgenstern axioms. On this measure, an alternative has higher expected utility just in case it would be preferred for the individual's sake after calm deliberation with all pertinent information while attending to this individual's self-interest only.⁷ Moreover, two alternatives have the same expected utility just in case such deliberation would yield indifference between them. If the individual is a normally capable adult with well-formed preferences, then the person expressing this preference can be taken to be an ideally rational version of this adult. If the individual is incapable of such deliberation (as we suppose Adam and Bill are), then the person expressing this preference can be taken to be an ideally rational guardian who has deliberated with exclusive focus on his charge's self-interest.

We will make one further assumption with regard to the interests that are taken into account in this measure of utility. We will assume that these reasons of self-interest do not include an interest in being fairly treated in the distributive process. Individual utilities therefore do not incorporate information on the degree to which the latter interest is satisfied.⁸ This assumption makes it possible to represent fairness considerations separately from well-being considerations.

Suppose that 1 is the utility of being fully sighted, 0.8 the utility of the considerable visual impairment, and 0.65 the utility of the severe impairment. The utilities associated with

⁷ This is not a proposal regarding what utility *consists in*. Rather, it is a proposal about how to *measure* the magnitude of a person's utility. One can maintain that utility consists of something other than preference satisfaction, while also maintaining that the specified idealized preferences faultlessly track the magnitude of this other thing.

⁸ In this respect, we depart from the assumption made in Fleurbaey (2010, 653-4).

treatments M_E and M_P in the two possible states of the world are shown in the top half of Table 1.

Table 1. The distribution of utility in two cases.

Alternative	Person	State of the world (equiprobable).	
		s_1	s_2
VISUAL IMPAIRMENT CASE 1			
M_E	Adam	0.8	0.8
	Bill	0.8	0.8
M_P	Adam	1	0.65
	Bill	0.65	1
VISUAL IMPAIRMENT CASE 2			
M_E	Adam	0.8	0.8
	Bill	0.8	0.8
M_U	Adam	1	1
	Bill	0.65	0.65

One of the aims of this chapter is to figure out how Teresa should decide VISUAL IMPAIRMENT CASE 1 under the assumption that she is an egalitarian, who rightly cares both about reducing outcome inequality and about increasing individuals' well-being.⁹ In cases of decision-making under *certainty*, because of the moral importance of improving the lot of those who are less well off than others, she has decisive reason to (and therefore does) prefer a medicine that ensures that both Adam and Bill end up with the considerable impairment to a medicine that leads one of them to retain excellent vision and leaves the other with the severe impairment.¹⁰ In other words, in a choice between M_E and M_U in the lower half of Table 1, Teresa would rightly choose M_E even though it has lower total utility, because for the utility values in question, the improvement in the condition of the worst off is worth the sacrifice in

⁹ Note, however, that our argument against *ex ante* Pareto is equally effective if one instead gives priority to the worse off because one is a Prioritarian, who gives additional weight to utility gains when they occur at a lower absolute level of utility. (For a discussion of Prioritarianism, see Parfit 2002.)

¹⁰ We do *not* assume that she regards *any* improvement in the condition in the worst off, no matter how small, to be more important than *any* improvement in total utility, no matter how large.

total utility. We assume, therefore, that such cases of certainty should not be resolved by an appeal to what would be agreed to by Adam and Bill's representatives if they were placed behind a so-called "hypothetical veil of ignorance" which assigns an equal probability to each of ending up in any person's shoes (see Harsanyi, 1953). (In this case, this would favour M_U , the alternative with the highest total utility.) As has been argued elsewhere, an appeal to such a device is an inappropriate way to make interpersonal tradeoffs, because it violates the separateness of persons and leads to what are, intuitively, unacceptably inequalitarian decisions.¹¹ One of the central questions of this chapter is therefore: given that Teresa ought to prefer M_E to M_U , how should Teresa decide between M_E and M_P ?

II. Against *ex ante* Pareto

On the measure of utility here assumed, it is true by definition that a risky alternative like M_P has higher expected utility for each person than a secure alternative like M_E just when, from the perspective of both Adam's and Bill's self-interest, the chance of the additional benefit offered by the risky alternative outweighs the risk of the foregone benefit associated with this alternative. If Teresa were to choose M_P over M_E , it may therefore seem that she could offer each of them the following justification for this choice:

In making this choice, I had to balance a 50 percent chance of your having excellent vision rather than the considerable visual impairment against a 50 percent chance of your having the severe visual impairment rather than the considerable visual impairment. I balanced this potential additional gain and potential gain forgone from

¹¹ See, for example, Nagel (1970, 138-9).

the perspective of your self-interest alone, in just the manner a guardian charged with looking after your interests, who had the same information I had, would have balanced them. On this way of balancing, the chance of the great benefit of unimpaired vision outweighed the associated risk of instead receiving the small benefit and thereby ending up with the severe impairment. In sum, by choosing M_P , I did the best I could for you, given the information I had at the time. I chose M_P for your sake.

However, this purported justification has a key difficulty.¹² It appeals to guardians who have the same information about their charges' possible fates as Teresa. (We are here imagining that Adam and Bill each have a guardian entirely to themselves, so that there are two guardians.) Such guardians are imperfect insofar as their information is limited. Because of this limited information, things that *seem* to such a guardian to be the best that can be done for his charge may not, in fact, be in his charge's interests. A guardian who is less often in error about what is, in fact, in his charge's interests would be a better guardian for an individual to have. Indeed, an ideal guardian would know all relevant information about what would, in fact, best serve his charge's interest. From Teresa's perspective, an informed guardian would also be ideal. The role of the guardians in the preceding purported justification is to represent the interests of Adam and Bill; and for each child, Teresa should want to know not merely which alternative seems to be in each child's interests, but also the extent to which each alternative will truly satisfy each child's interests.

The following considerations further underline that the justificatory force of the assent of uninformed guardians must be treated with suspicion. Suppose that Teresa must choose

¹² A similar objection to appealing to people's imperfectly informed preferences is offered in Fleurbaey (2010, 659).

between M_E and M_U —she therefore knows the impact each treatment will have on each named individual. Suppose further that she is confronted with imperfectly informed guardians. Because of their lack of information, each guardian reasonably regards M_U as involving a fifty-fifty gamble between a full cure and the severe impairment for his charge (in just the manner in which M_P is regarded by the imperfectly informed guardians imagined in the purported justification for M_P). Each guardian therefore reasonably (given his beliefs) favours M_U for his charge. In these circumstances, Teresa obviously could not justify the choice of M_U to Bill on the grounds that his uninformed guardian regarded it as in his interest. After all, she knows that his guardian’s preference for M_U over M_E is merely due to ignorance.

Now suppose that Teresa must choose between M_E and M_U , but that because of lack of information, she, Adam’s guardian, and Bill’s guardian all regard M_U as involving a fifty-fifty gamble between a full cure and the severe impairment for each child. Suppose she can easily and costlessly acquire all relevant information about the impact of M_U . In these circumstances, Teresa obviously should acquire this information. Moreover, if she fails to do so, she cannot justify the choice of M_U to Bill on the grounds that she and his uninformed guardian regarded it as in his interest.

We conclude that whenever possible, Teresa should base her decisions on what *fully-informed* guardians, who know their charges’ true interests, would tell her about the extent to which alternatives are in their charges’ interests. Now, it may seem that this conclusion does not undermine the appeal to the assent of uninformed guardians in VISUAL IMPAIRMENT CASE 1. For unlike in the two cases just mentioned, in this case, Teresa cannot know what Adam’s (or Bill’s) fully-informed guardian would advise her. However, she *can* infer some of the things about the preferences of their fully-informed guardians. For one, she can infer that there would be no unanimity among such fully-informed guardians in favour of M_P . One of them would favour M_E on behalf of his charge; the other would favour M_P .

Teresa can also infer that if she consulted such fully-informed guardians, insofar as she bases her decision on individuals' well-being, she would rightly no longer regard it as justifiable to each to provide both with M_P . For if these fully-informed guardians knew that the state of the world would turn out to be s_1 , and therefore told her that M_P was in Adam's interest but not in Bill's, then, insofar as she bases her decision on these interests, she would regard the choice of M_P as unjustifiable to Bill. She would not regard it as justifiable to Bill because, in line with the assumptions made in Section I, it is unjustifiable to sacrifice the interests of someone who, as things turn out, will be no better off than the other person under either policy in order to make that other person better off still, at least for the size of the benefits under consideration. Moreover, if these fully-informed guardians knew that the state of the world would turn out to be s_2 , and therefore told Teresa that M_P was in Bill's interest but not in Adam's, then she would, for the same type of reason, regard the choice of M_P as unjustifiable to Adam.

Teresa also knows that, insofar as she bases her decision on how the available medicines will affect individuals' true well-being, she should favour M_E over M_P no matter what the fully-informed guardians would tell her about their charges' interests. If these fully-informed guardians knew that the state of the world would turn out to be s_1 , and therefore told her that M_P was in Adam's interest but not in Bill's, then, in line with the extra weight she should assign to helping the worst off, she should resolve these conflicting interests in Bill's favour, and choose M_E . And if these fully-informed guardians knew that the state of the world would turn out to be s_2 , and therefore told her that M_P was in Bill's interest but not in Adam's, then she should resolve these conflicting interests in Adam's favour, and again choose M_E .

In sum, Teresa can infer that if she could consult both Adam's and Bill's fully-informed guardians (a) M_P would not be justifiable to each; and (b) she should always favour policy M_E . (As mentioned, both (a) and (b) hold insofar as she bases her assessment on the

final distribution of well-being among Adam and Bill.) Now, to consult such fully-informed guardians is just to possess full relevant information. Because she cares about the real impact of alternatives on Adam and Bill's well-being, Teresa regards judgments made with full information about the well-being interests involved as superior to judgments made without full information. And in this case, once she has gone through the foregoing reasoning, Teresa knows the alternative that she should strictly prefer if she had full information, no matter what that information turns out to be. It seems a basic principle of rationality that she should decide in line with this preference (Fleurbaey, 2010; Adler, 2011).

The principle of rationality here invoked can be put as follows:

PRINCIPLE OF FULL INFORMATION, PART 1: When one knows that, in every state of the world with positive probability, one would rightly rank two alternatives in a particular way, then one should so rank them.¹³

¹³ This is weaker than the principle, common in decision theory, that if the set of feasible alternatives is unaffected by the acquisition of information, and one would not make any errors in deliberating on the basis of this information, it is always good or indifferent to acquire more information (Savage 1972). The former only pertains to cases where one knows what one would decide with full information, whereas the latter would tell one to always decide as one would with *more* information than one has now.

Note also that in applying the principle, one must take care when the outcome of a prospect depends on the absence of information at some earlier time. For example, suppose your alternatives are "go for a run" and "watch the Oxford-Cambridge boat race", that the possible states of the world are "Oxford wins" and "Cambridge wins", and that you do not know who will win. Suppose further that your ignorance of the outcome of the race does not affect the value of going for a run, but is crucial to your enjoyment when watching the race. In applying the PRINCIPLE OF FULL INFORMATION, you must compare the outcome of each alternative in these two states of the world. Since your ignorance of who will win affects your experience of watching the race, this does *not* involve comparing the outcome of "go for a run" with the outcome of "watch the race knowing that Oxford will win" and "go for a run" with "watch the race knowing that Cambridge will win". Rather, it involves

In sum, we have argued that in VISUAL IMPAIRMENT CASE 1, *ex ante* Pareto violates this principle of rationality, and should therefore be rejected in this case.¹⁴

It is important to note that the scope of this argument is limited. For this argument has no bite when one cannot infer *ex ante* that, with full information, one would always regard the same alternative as best. It therefore implies nothing about cases involving pure intrapersonal tradeoffs in which only one person's interests are at stake, or in which everyone's well-being outcomes are perfectly correlated. By way of illustration, consider the following ONE-PERSON CASE, which differs from our previous cases in that only Adam's

comparing the outcome "go for a run", with the outcome "watch the race in ignorance and Oxford wins", etc. The PRINCIPLE OF FULL INFORMATION then requires that if, in each state of the world, the outcome of watching in ignorance is better than the outcome of going for a run, then you should prefer watching the race.

¹⁴ This argument directly questions the soundness of the *ex ante* Pareto principle. An indirect argument is also available by reference to Harsanyi's Aggregation Theorem (Harsanyi 1955; for a non-technical presentation and detailed discussion of its implications, see Broome 1991). According to this theorem, the *ex ante* Pareto principle, combined with the assumption that social evaluation is made by a criterion that takes the form of the expected value of a social welfare function, implies that the social welfare function must be affine in individual utilities. Affine functions are such that the weight of a particular individual's utility is fixed and does not vary with how well off the individual is. If one assumes, as we have done here, that some priority must be given to the worse off, such affine functions are unacceptable. If, in addition, one assumes that social evaluation must be an expected value criterion, then Harsanyi's theorem implies that *ex ante* Pareto can no longer be satisfied.

Here, we did not make the expected value assumption about Teresa's evaluations, but only assumed that if Teresa would make the same decision in every state of the world, she should make it *ex ante*. This dominance principle is logically weaker than the expected utility hypothesis, but is enough to imply a rejection of *ex ante* Pareto. (A version of Harsanyi's theorem in which the expected social value hypothesis is replaced by the dominance principle is given in Fleurbaey 2009).

interests are at stake. She must choose between the certain option M_E ,¹⁵ which will ensure Adam ends up 0.8, and M_P , which will give Adam a one in two chance of ending up at 1—if s_1 turns out to be the state of the world—and a one in two chance of ending up at 0.65—if s_2 turns out to be the state of the world. In this case, if Teresa knew that s_1 will turn out to be the case, then she would prefer M_P ; if she knew that s_2 will turn out to be the case, then she would prefer M_E . In this case, therefore, Teresa cannot now ensure that her choice will be consistent with how she would choose with full information. The PRINCIPLE OF FULL INFORMATION therefore says nothing in this case, and our argument does not challenge the application of the *ex ante* Pareto principle to such cases.¹⁶

III. Extending the Principle of Full Information

In VISUAL IMPAIRMENT CASE 1, Teresa can infer that she would always strictly prefer the same alternative with full knowledge. It is natural to extend the PRINCIPLE OF FULL INFORMATION to cases in which Teresa is not certain which of two alternatives she would strictly prefer with full knowledge, but correctly infers that, with full knowledge, she might disprefer one of the two alternatives but would never disprefer the other. In this case, it seems straightforward that she should avoid the risk of choosing an alternative she would disprefer with full knowledge, and choose the alternative she knows she will never disprefer with full

¹⁵ Note that, insofar as Adam's well-being is considered in isolation from how well or badly off others are, egalitarianism has nothing to say on the choice between M_E and M_P . In this sense, and unlike in our two-person cases, M_E is not the egalitarian option in this case.

¹⁶ Indeed, elsewhere, we argue that in such cases, insofar as she considers Adam's well-being in isolation, it is reasonable for Teresa to do what the *ex ante* Pareto principle recommends. See Voorhoeve and Fleurbaey (forthcoming) and Otsuka and Voorhoeve (2009).

knowledge. She can thereby ensure her choice is consistent with her perfectly informed preference. This reasoning yields:

PRINCIPLE OF FULL INFORMATION, PART 2: When one knows that there is only one way of ranking two alternatives that will ensure one will never choose an alternative that one would rightly disprefer in full knowledge of the state of the world, then one should so rank them.

By way of illustration, consider the case described in Table 2 (in which M_M might, but need not, generate outcome inequality). Insofar as she is concerned with well-being, Teresa can infer that, if she were to learn that s_1 will be the case, then she would prefer M_E to M_M , while if she were to learn that s_2 will be the case, then she would be indifferent between these two medicines. Once she has made this inference, she knows that she can certainly avoid choosing contrary to her fully-informed preference by choosing M_E . It seems obvious that she should so choose.

Table 2. The distribution of utility in one case.

Alternative	Person	State of the world (equiprobable).	
		s_1	s_2
VISUAL IMPAIRMENT CASE 3			
M_E	Adam	0.8	0.8
	Bill	0.8	0.8
M_M	Adam	1	0.8
	Bill	0.65	0.8

A final extension concerns cases in which Teresa can infer that, with full information, she would always be indifferent between two alternatives. In this case, it seems she should regard both as equally good. This yields:

PRINCIPLE OF FULL INFORMATION, PART 3: When one knows that, in every state of the world with positive probability, one is indifferent between two alternatives, then one should be indifferent between these alternatives.

By way of illustration, suppose that Teresa must choose between M_P and M_U alone, as in Table 3.¹⁷

Table 3. The distribution of utility in one case.

Alternative	Person	State of the world (equiprobable).	
		s_1	s_2
VISUAL IMPAIRMENT CASE 4			
M_P	Adam	1	0.65
	Bill	0.65	1
M_U	Adam	1	1
	Bill	0.65	0.65

Insofar as Teresa is concerned with final well-being alone, she is indifferent between the outcomes of these medicines in both s_1 and in s_2 . If well-being outcomes are all that is relevant, then she should therefore be indifferent between these medicines.

So far, we have considered only one part of the evaluation of alternatives, namely, their impact on individuals' well-being interests. But other things besides well-being interests might matter. One might think, for instance, that M_P is less unfair than M_U , because M_P , but not M_U , offers Bill a chance of the benefit of unimpaired vision. We will now examine such fairness claims and show how an appeal to the PRINCIPLE OF FULL INFORMATION is consistent with giving them due weight.

¹⁷ In this case, our mnemonic names for the medicines are less apt, since both M_P and M_U are *ex ante* Pareto optimal and both maximize expected utility. However, we continue to use them because the effects of these medicines are the same as in our central case.

IV. Fairness and the Principle of Full Information

Broadly speaking, there are two leading explanations of the idea that an alternative such as M_P , which gives each person an equal chance at a great benefit that only one can receive, is fairer than an alternative such as M_U , in which it is known at the time of choice who will receive the great benefit. Following Wasserman (1996), we will refer to these as the “prophylactic view” and the “distributive view.”

The prophylactic view holds that giving each an equal chance of this benefit is fairer when and because it ensures the absence of partiality or favouritism in the distribution of a good. We will assume that Teresa is publicly known to be fair in this sense. Teresa therefore has no prophylactic reasons for favouring alternatives that give each person an equal chance of ending up in the best position; nor does she have reasons to choose such an alternative to express or demonstrate her impartiality, unless, of course, impartial consideration of the interests of each already favours such an alternative.

On the distributive view, a given outcome inequality among people with equally strong claims to a benefit is less unfair when each person has a chance to end up better off than when the worse off have no such chance, because in receiving this chance, each person receives an equal share of something of expected value.¹⁸ On this view, the more equally this expected value is distributed at a relevant point in time, the fairer the distribution. Of course, the *well-being value* of a chance for a person evanesces once it is clear that this chance is

¹⁸ This view’s judgment that a given pattern of outcome inequality is less unfair when individuals have equal chances of ending up in any position is widely shared. See, for example, Diamond (1967), Broome (1990), and Otsuka (ms.).

unrealized.¹⁹ In this sense, having a chance at a benefit which does not materialize neither compensates a person for the well-being foregone nor offers any satisfaction of his well-being interests. However, a chance's contribution to fairness does *not* evanesce—it remains true of someone who received an equal chance at an indivisible benefit of a given size, but for whom the benefit did not materialize, that he had a fair chance of receiving it. For this reason, this equal chance contributes to the satisfaction of each individual's interest in being treated fairly in the distributive process, where this fair treatment involves more than just having his interests given equal consideration.²⁰

Now, some, including Wasserman (1996), reject the distributive view and with it, the idea that equal chances make a contribution to fairness beyond their contribution to ensuring and expressing impartiality. We will not here pronounce on whether the distributive view is correct. Instead, we explore how Teresa should choose if only the prophylactic view is correct and how she should choose if, instead, the distributive view is correct.

Let us start with VISUAL IMPAIRMENT CASE 4. If only the prophylactic view is correct, then all that matters is individuals' well-being, so that Teresa is rationally required to be indifferent between M_P and M_U . By contrast, if the distributive view is correct, then the well-being information listed in Table 3 is not all that is relevant to the evaluation of these alternatives. When considering the outcomes of these alternatives in each state of the world, we need to consider *both* each individual's well-being *and* the chance he had to achieve a

¹⁹ In our cases, it is assumed that individuals' chances of receiving a good are non-transferable, so that they have no opportunity to exchange their as-yet-unrealized chance for some other good that will contribute to their well-being. We can also assume that the individuals concerned do not know their chances of receiving a benefit, so that they derive no benefit from the mere anticipation of the possible gain. The chance of receiving a benefit therefore makes no contribution to the quality of life of a person who does not end up receiving it.

²⁰ As mentioned in Section I, we assume that the satisfaction of this interest does not figure in our measure of utility.

different outcome (Broome 1991, 111-15). The relevant outcomes for the choice between M_P and M_U are then those depicted in Table 4.

Table 4. The outcomes in one case if the distributive view is correct.

Alternative	Person	State of the world (equiprobable).	
		s_1	s_2
VISUAL IMPAIRMENT CASE 4			
M_P	Adam	1 & had a one in two chance to be severely impaired.	0.65 & had a one in two chance to be fully sighted.
	Bill	0.65 & had a one in two chance to be fully sighted.	1 & had a one in two chance to be severely impaired.
M_U	Adam	1 & had no chance of any other outcome.	1 & had no chance of any other outcome.
	Bill	0.65 & had no chance of any other outcome.	0.65 & had no chance of any other outcome.

When the consequences are specified in this manner, the PRINCIPLE OF FULL INFORMATION requires that M_P be chosen in this case, as it is preferable to (because less unfair than) M_U in every state of the world.

Let us now return to VISUAL IMPAIRMENT CASE 1. If only the prophylactic view is correct, then the reasoning of Section II applies, and Teresa should choose the egalitarian option. If the distributive view is correct, then the relevant outcomes are depicted in Table 5.

Table 5. The relevant outcomes in one case if the distributive view is correct.

Alternative	Person	State of the world (equiprobable).	
		s_1	s_2
VISUAL IMPAIRMENT CASE 1			
M_E	Adam	0.8 & had no chance of any other outcome.	0.8 & had no chance of any other outcome.
	Bill	0.8 & had no chance of any other outcome.	0.8 & had no chance of any other outcome.
M_P	Adam	1 & had a one in two chance to be severely impaired	0.65 & had a one in two chance to be fully sighted,.
	Bill	0.65 & had a one in two chance to be fully sighted,.	1 & had a one in two chance to be severely impaired.

If the distributive view is correct, then the PRINCIPLE OF FULL INFORMATION requires that Teresa's preference should entirely track her preference over the following outcomes: (1) Adam and Bill both have the considerable impairment and never had a chance of anything better; (2) one of them is fully sighted, and one of them has the severe impairment, with the unfairness of this outcome being mitigated by the fact that the severely impaired person had a one in two chance of instead ending up unimpaired which is equal to the chance that the unimpaired person had of instead ending up with the severe impairment.

Now, the choice between (1) and (2) depends on how much weight Teresa should give to improving the lot of those who end up worst off, and how significant it is that the unfairness of receiving the lesser benefit is mitigated by the fact that one could have ended up with the greater benefit. If Teresa should give great weight to improving the lot of those who are less well off than others and if equal chances do little to mitigate the unfairness of unequal outcomes, then she should prefer M_E over M_P . If, by contrast, Teresa should not sacrifice much total utility for the sake of the worst off, then the fact that the equal chances offered by M_P mitigate the associated outcome-unfairness may tip the balance in its favour. Whatever the right way to balance these concerns, it is important to note that Teresa can

ensure that the alternative she chooses is one she would regard as best with full knowledge of the relevant outcomes.²¹

We can conclude that so long as Teresa knows which view of fairness is correct, in VISUAL IMPAIRMENT CASE 1, she can make a decision that she can be sure will be in line with the preferences she would have with full information. If she knows that only the prophylactic view is correct, then she can infer that, with full information, she would always choose the egalitarian alternative. If, by contrast, she knows the distributive view is correct, then, in order to arrive at a decision she would approve of with full information, she should ask herself: “Is an outcome in which both people are considerably impaired at least as good as the outcome in which one person ends up fully sighted and the other severely impaired and this severely impaired person had a fair, equal chance at a full cure?” and choose the egalitarian option just in case the answer is “yes.”

V. Return to Real-World Cases

Our discussion of our visual impairment cases has implications for closely analogous real-world cases. As relatively wealthy private individuals, most of us in the developed world may find ourselves in Teresa’s position when we consider how to allocate resources which can do good for the children of strangers. If our argument is correct, then *ex ante* Pareto should play no role in our deliberations about such cases.

²¹ At least, so long as she also knows the nature of the chances involved (e.g. whether they are merely subjective, objective, or indeterministic). This caveat is required because one might argue that objective or indeterministic chances make a greater contribution to fairness than merely subjective chances (Otsuka, ms.).

Our conclusions about real-world cases that are less like our stylized case must be more speculative. With this caveat in mind, what can we say about a case in which Adam and Bill are not children, but rational adults who have the same information that Teresa has? Adam and Bill will then prefer that Teresa administers the *ex ante* Pareto-optimal medicine. One may think that, since they are rational adults, Teresa has reason to defer to their judgment in order to promote their autonomy. We would question this defence of *ex ante* Pareto on two grounds. First, Adam and Bill's judgment is imperfectly informed. Even if deferring to their *fully*-informed preferences promotes their autonomy, it is not clear that deferring to their *imperfectly*-informed preferences promotes their autonomy (Harsanyi 1982).

Second, a plausible understanding of individual autonomy requires only that one defer to a person's wishes regarding choices he has a right to make. On this understanding, neither Adam's nor Bill's autonomy would be threatened if Teresa acted contrary to their wishes by giving them the egalitarian treatment. Since, by hypothesis, she rightfully controls the treatments, if she were to give them the egalitarian treatment, Teresa would not interfere with any choices they have a right to make.²²

In sum, even if Adam and Bill are adults, is doubtful that Teresa has an autonomy-based reason to do what they want. We tentatively conclude that even in cases where adults' well-being is at stake, someone in Teresa's position should regard *ex ante* Pareto as having no force.

What if the decision-maker is not a stranger, but is instead Adam and Bill's agent, charged by each of them with doing what is in his expected interest? In this case, such a decision-maker may well have good reason to do what both prefer her to do. This is relevant for the MAMMOGRAM CASE, insofar as the officials on the Task Force and the doctors who would be acting on its recommendations are such agents of the women affected. Insofar as

²² See Otsuka and Voorhoeve (2009, p. 186).

they are agents of this kind, in this case, it would be inappropriate for them to advise women to do something which is not expectedly beneficial for these women.

Nonetheless, it seems to us that the members of the Task Force are not *merely* agents charged by these women with promoting their *ex ante* prudential interests. As public officials, they are also tasked with ensuring distributive justice in health-related quality of life. As such, we take it that, in cases of certainty, they should give special weight to improving the condition of those who end up worse off than others. Moreover, in cases of risk, they should be interested in what is truly in each woman's well-being interests, and not merely in what, due to lack of information, seems to be in these women's interests. Members of the Task Force should therefore not regard the fact that an alternative is marginally expectedly beneficial for each as a sufficient justification for choosing this alternative.

By way of illustration, imagine that the feasible set of the Task Force is expanded to include an alternative which we will call REALLOCATION AND ROUTINE SCREENING. This involves the use of medical resources in order to somewhat lower the expected costs of routine mammography to the point where undergoing it is *not* a net expected burden to women. (This might be done, say, by investing in training or equipment which lowers the incidence of false positive test results.) If this alternative were implemented, it would be unproblematic for doctors to advise their patients to undergo routine screening. Suppose that these resources would be reallocated from services which provided modest benefits to each woman, and that, as a consequence, this alternative is *ex ante* Pareto-inferior to NO ROUTINE SCREENING, which involves no such reallocation. However, because it prevents many early deaths, from the perspective of the distribution of the satisfaction of women's genuine well-being interests, the outcome of REALLOCATION AND ROUTINE SCREENING might be better than the outcome of NO ROUTINE SCREENING. If so, then the members of the Task Force would have reason to choose the former, contrary to *ex ante* Pareto.

Conclusion

We have reviewed cases in which an alternative that is in the expected interests of each will leave those who end up worst off substantially less well off. We have argued that in such cases, we should not heed the call of that famous Italian siren, *ex ante* Pareto. We should not base our decisions on what, due to a lack of information about the impact of our choices on each individual, merely seems to be in every individual's interest. Instead, we should ask what we would decide with full knowledge about the state of the world, and choose accordingly.

References

- Adler, Matthew. *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. (Oxford: Oxford University Press, 2011).
- Broome, John. "Fairness," *Proceedings of the Aristotelian Society, New Series* 91 (1990): 87-101.
- Broome, John. *Weighing Goods* (Oxford: Blackwell, 1991).
- Diamond, Peter. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment," *Journal of Political Economy*, 75 (1967): 765-6.
- Fleurbaey, Marc. "Two Variants of Harsanyi's Aggregation Theorem," *Economics Letters* 105 (2009): 300-2.

- Fleurbaey, Marc. "Assessing Risky Social Situations," *Journal of Political Economy* 118 (2010): 649-80.
- Harsanyi, John C. "Cardinal Utility in Welfare Economics and the Theory of Risk-Taking," *Journal of Political Economy* 61 (1953): 434-5.
- Harsanyi, John C. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy* 63 (1955): 309-2.
- Harsanyi, John C. "Morality and the Theory of Rational Behaviour," in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press, 1982), pp. 39-63.
- Nagel, Thomas. *The Possibility of Altruism*. Princeton: Princeton University Press, 1970.
- Otsuka, Michael. "The Fairness of Equal Chances," (ms.).
- Otsuka, Michael and Alex Voorhoeve. "Why It Matters that Some Are Worse Off than Others," *Philosophy and Public Affairs* 37 (2009): 171-99.
- Parfit, Derek. "Equality or Priority?" in *The Ideal of Equality*, ed. Matthew Clayton and Andrew Williams (Basingstoke: Palgrave, 2002), pp. 81–125.
- Savage, Leonard. *The Foundations of Statistics*, 2nd revised ed. (New York: Dover Publications, 1972).
- U.S. Preventive Services Task Force. "Screening for Breast Cancer: U.S. Preventive Services Task Force Recommendation Statement," *Annals of Internal Medicine* 151 (2009): 716-26.
- Voorhoeve, Alex and Marc Fleurbaey. "Egalitarianism and the Separateness of Persons," *Utilitas* (forthcoming).
- Wasserman, David. "Let Them Eat Chances: Probability and Distributive Justice," *Economics and Philosophy* 12 (1996): 29-49.