

The Ethics of Making Risky Decisions for Others

Luc Bovens
UNC at Chapel Hill
Department of Philosophy
LBovens@UNC.edu

1. Introduction

Utilitarianism, it is often said, is not sensitive to the distribution of welfare across different individuals'.¹ In a risk-free environment, what this means is that utilitarianism is indifferent whether an increment in welfare of a given size goes to a better-off or a less well-off person, and it requires an intervention that procures even a slightly greater increase in welfare to the better off at the cost of a smaller decrease in welfare to the worse off.

One can construe this as a kind of dilemma. On the one hand, the utilitarian has a reasonable objective, that is, to maximise total or average welfare in society. On the other hand, the sensitivity of their opponent is reasonable as well: The worse off should not be expected to incur losses for the sake of just marginally larger gains to the better off. Different decision-makers will resolve this dilemma in different ways depending on how sensitive they are to the distribution of welfare.

In the presence of risk, the matter is much more complex than in a risk-free environment because there are more distributions in play that elicit distinct sensitivities. These sensitivities will generate different types of dilemmas.

In section 2, I will present four dilemmas that all involve risk in decision-making for others. These dilemmas seem quite disparate. In section 3, I construct a model that captures four distinct distributional sensitivities that a decision-maker might have when making risky decisions for others. The model specifies four parameters that quantify the strength of these sensitivities. In section 4, I will show how each of the four dilemmas occurs because a particular distributional sensitivity kicks in and pulls us away from what the utilitarian would prescribe in the context of decision-making under risk. In section 5, I will turn to the much-discussed case of *Saving the Few at Greater Risk to the Many*. This case is more complex, because there are two distributional sensitivities at work that pull in opposite directions from the utilitarian calculus. Finally, in section 6, I will consider objections to each of the four distributional sensitivities. Section 7 reflects on the role of formal modelling by means of prospects.

2. Four Dilemmas

¹ For example, see Gauthier 1963: 121–7, Rawls 1974–75, Sen 1973, and Williams 1973.

2.1 Restrictive Intervention

Suppose that we have a cheap but less effective drug that we can provide to all, or an expensive but more effective drug that, due to budgetary constraints, we can provide only to half the population. The catch is that the expected number of lives saved by the more expensive drug programme, even if only provided to half, is greater than the expected number of lives saved by the cheaper drug programme, provided to all. Should we procure the cheaper drug to all or the more expensive drug to half?

Ubel *et al.* 1996 construct a case in which the allocation of the expensive drug is conducted by randomisation. They polled decision theorists, prospective jurors and medical ethicists asking whether we should procure the cheaper drug to all or the more expensive drug to half. This case constitutes a dilemma with forces pulling in opposite directions. On the one hand, one wants to maximise the expected number of people saved. This favours the more expensive drug. On the other hand, once the randomisation device has determined who will get the more expensive drug, people have unequal chances of survival. This favours the cheaper drug. Indeed, Ubel *et al.* find that decision theorists tend to favour the more expensive drug for half, while prospective jurors and medical ethicists tend to favour the cheaper drug for all.

2.2 Diversification

The *GiveWell* website, run by Effective Altruists, recommends charities that provide the greatest expected reduction in suffering and premature death for each marginal dollar contributed. Organisations that do extremely well, according to this calculus, are anti-malaria and deworming initiatives.

Critics ask why *GiveWell's* calculus is not more favourable to otherwise highly-regarded charities like Oxfam, MSF (*Médecins Sans Frontières*) and the Red Cross. The response from Effective Altruists is that they respect these organisations, but that organisations that support diverse causes do worse in their calculus than organisations that support single causes. One reason for this is that a diversified portfolio tends to include causes with effects that are hard to measure, such as advocacy against corruption.

However, there is another difference between single-cause and multiple-cause organisations. For single-cause organisations, the chances that your marginal dollar will do some good are more correlated than for multiple-cause organisations.

Consider deworming initiatives. These initiatives support massive deworming in schools in the developing world. The pupils are not tested individually, since treatment is cheaper than testing and the pills are relatively safe. The original study by Miguel and Kremer 2003 suggested that deworming has huge educational impacts (which hold the promise of future economic benefits). Taylor-Robinson *et al.* 2015, Aiken *et al.* 2015, and Davey *et al.* 2015 cast doubt on these results. However, *GiveWell* continues to support the initiatives, because they

are cost effective. Deworming is cheap and it just *might* have a huge impact. However, a blogger on the *GiveWell* website also admits that it “might have close to zero impact” (Sean 2017).

Granted, deworming may not work equally well in all places, but one would expect that the success or failure of these interventions around the developing world is at least to some extent correlated. Hence, in giving to deworming initiatives, there is a non-negligible chance that one’s donations will have little or no impact. By contrast, contributions to organisations like Oxfam go to a broad array of causes. The expected reduction in suffering and premature death from donations may be lower than from deworming initiatives, but at least a portion of one’s donation is likely to do some good.

The dilemma is this: Should we make charitable donations to single-cause organisations in order to maximise the expected benefits, even if there is a good chance that the money will be money down the drain? Or, should we donate to multiple-cause organisations in order to make sure that our donations will at least do some good in the world, even if the expected benefits of our donations are lower?

2.3 Unwelcome Risk Reductions

Spiegelhalter 2015 has constructed a flow graph that offers a visual representation of the costs and benefits of breast cancer screening in the UK. The graph shows that 7.5% of women between 50 and 70 in the UK have symptoms of breast cancer that would be discovered by screening. If they are screened, then all 7.5% will be treated, 6% will survive and 1.5% will die. If they are not screened, then 1.5% will have symptoms that would be discovered in screening, but they will never become aware of it: They will remain undiagnosed, untreated and the disease will never bother them. In other words, if they had been screened, this 1.5% would have been needlessly treated. The remaining 6% will eventually show symptoms and be treated; 4% will survive and 2% will die.

So, screening reduces the chance of mortality with $2\% - 1.5\% = .5\%$. But it comes at a cost of 1.5% who are needlessly treated. If that 1.5% would not have been screened, they would have never known that they had symptoms of breast cancer.

Many women opt against screening when told about these stats. A mere .5% reduction in mortality does not justify the hassle for them. There is the annoyance of the screening, the needless agony of false positives, and the 1.5% of treatments that are unnecessary. And yet, screening does save about 1300 lives per year in the UK.

This is our dilemma. On the one hand, it seems that we should respect the wishes of fully informed patients. On the other hand, there is a compelling case to be made that, taking into account costs, saving an expected 1300 lives merits the implementation of a screening process.

2.4 Not on My Watch

Johnson and Rehavi 2016 present a study of the percentage of deliveries by C-sections for prospective mothers who are themselves MDs versus for prospective mothers who are not MDs. They find that the percentage of unscheduled C-sections—that is after a vaginal birth was attempted—is higher for prospective mothers who are not MDs.

Why is this the case? It may be the case that some obstetricians need to make their car payments by the end of the month and perform needless C-sections. C-sections are more lucrative than vaginal births and under conditions of asymmetry of information, financial incentives create physician induced demand for C-sections. However, they can't pull this off on prospective mothers who are MDs, since those patients know better. One datum in support of this hypothesis is that we also find a lower rate of C-sections on partners of prospective fathers who are MDs and that the rate of C-sections is even lower when we focus on prospective mothers who are obstetricians. The results do not hold for scheduled C-sections, i.e., C-sections that were planned beforehand. The authors suggest that this is the case because there are strict criteria for scheduled C-sections.

Another explanation is that the obstetricians are afraid of being sued by the women or their families for *not* proceeding with a C-section during somewhat difficult or risky deliveries. They may trust that prospective mothers or partners of prospective fathers who are themselves MDs are less likely to file suit against a fellow doctor for medical malpractice.

However, let us try to be less cynical and attempt a friendlier explanation. When there are concerns during the birthing process, proceeding to a C-section is often the less risky option. Vaginal birth can result in success (a healthy mother and foetus) or failure (death or injury to the mother or foetus). Due to the longer recovery period, the outcome of a C-section typically ranks between an outcome of success and an outcome of failure of a vaginal birth.

If the prospective mother is not an MD, then the obstetrician takes the decision *for* the prospective mother. Even if holding out for a vaginal birth is a risk that the obstetrician themselves would want to take if they were in the position of the prospective mother, they may not want to take this risk *for* the prospective mother. Hence, they proceed to a C-section more swiftly.

To defend this decision, they may say: "I can't have death or injury due to forgoing a C-section happen *on my watch*." And this is not (just) because they are concerned about their reputation or about being sued. Rather, a good person finds it harder to take a loss when gambling with the welfare of others, than with the welfare of themselves. This is why they might want to take the gamble when watching over themselves, but not when watching over others.

If the prospective mother is herself an MD, then the obstetrician is more like a fiduciary caretaker. They simply adopt the risk attitudes of the MD prospective mother. The obstetrician doesn't take the decision *for* her – rather they see themselves as executing the decision that

the prospective mother would take for herself. Hence, they are less swift in proceeding to a C-section.²

The dilemma can then be phrased as follows. The presumption is that non-MD prospective mothers, if they were fully informed, would have risk attitudes that are no different from MD prospective mothers. Should an obstetrician respect the risk attitudes that the non-MD prospective mother would have if she were fully informed? Or, is it permissible for the obstetrician to import additional risk aversion over and above the risk attitudes of the non-MD prospective mother?

3. Distributional Sensitivities

I will construct a model to cast light on where the pull comes from in each of these dilemmas.

3.1 Four Distributional Sensitivities

Suppose that you were asked to make a choice between prospects *Risky* and *Certain* in Tables 1 and 2 that will affect two people. A matrix represents a *prospect*, rows represent *personal prospects*, columns represent *states* (of the world). There is a probability distribution defined over the states and the entries in the matrix are utility values.

² Granted, one problem with this explanation is that the difference in C-section rates performed on non-MDs versus MDs disappears when we restrict ourselves to HMO hospitals in which obstetricians get a fixed salary and don't work on a fee-for-service basis. This favours the hypothesis of physician-induced demand driven by financial incentives.

Expected Utility (EU)		S1 <i>Prob(S1) = .3</i>	S2 <i>Prob(S1) = .7</i>
6.7	Person 1	20	1
3.4	Person 2	2	4
<i>Average EU Expectation of SA 5.05</i>	State Average (SA)	11	2.5

Table 1. *Risky*

Expected Utility		S <i>Prob(S) = 1</i>
5	Person 1	5
5	Person 2	5
<i>Average EU Expectation of SA 5</i>	State Average (SA)	5

Table 2. *Certain*

We can calculate the *value* of the prospect *ex ante*, that is, by computing the average of the expected utilities. Or, we can calculate the value of the prospect *ex post*, that is, by computing the expectation of the average utility in each state. By simple algebra, the average of the expectations equals the expectation of the state averages.

Harsanyi 1955 imposes a set of requirements of minimal rationality on the ranking of prospects and shows, in his aggregation theorem, that we should rank prospects relative to the *weighted* average of the expectations in the prospect. If we add to this a requirement of anonymity—that is, the identity of the people does not matter—then we are required to order prospects relative to the average of the expectations of the persons in the prospect. Hence, following Harsanyi, we should rank *Risky* over *Certain*.

But forget about theory for a second. Suppose that a decision-maker is entrusted with making a decision between these two prospects. I submit that many decision-makers would want to play it safe and choose *Certain* over *Risky*. What might such a decision-maker say if they were asked to give a justification for their decision? They might mention a number of things that they dislike about *Risky*:

- (i) They may dislike the fact that person 2's expectation is substantially lower than person 1's.
- (ii) They may dislike the fact that state S2, which is likely to materialise, has a low average utility.
- (iii) They may dislike the fact that whatever state materialises, there will be inequality between the persons.
- (iv) They may dislike taking the risk of receiving a low outcome on behalf of person 1 if S2 materialises or a low outcome on behalf of person 2 if S1 materialises.

The decision-maker finds fault with four types of distributions: In (i), they find fault with the *inter-personal prospect distribution*. This is the distribution of the expected utilities, viz. $\langle 6.7, 3.4 \rangle$. In (ii), they find fault with the *inter-state distribution*. This is the distribution of the state averages, viz. $\langle 11, 2.5 \rangle$. In (iii), they find fault with the *intra-state distributions*. There are two intra-state distributions, viz. $\langle 20, 2 \rangle$ and $\langle 1, 4 \rangle$. In (iv), they find fault with the *intra-personal prospect distributions*. There are two intra-personal prospect distributions, viz. $\langle 20, 1 \rangle$ and $\langle 2, 4 \rangle$.

The decision-maker resists the recommendation of the prospect with the greatest average expectation or the greatest expectation of averages by pointing to four worrisome distributional features. They are sensitive to spread, that is, they would rather see less than more spread in these distributions and this is what motivates the decision-maker to recommend B over A. For simplicity, let us name the average of the expectations or the expectation of the state averages the *utilitarian* value of the prospect, though the import of Harsanyi's Theorem is open to interpretation (see Weymark 1991 and Greaves 2017). Hence, the utilitarian value of *Risky* may exceed the utilitarian value of *Certain* and yet, we would prefer *Certain* over *Risky* on distributional grounds.

3.2 Modelling Distributional Sensitivities

I will import these distributional sensitivities in the determination of the value of a prospect.

3.2.1 The Inter-Personal Prospect Distribution

Suppose that we are sensitive to the inter-personal prospect distribution. Let us take some simple numbers. Suppose that person 1 has an expectation of 16 and person 2 has an expectation of 4, so that the utilitarian value of the prospect is 10.

A decision-maker who would rather see less of a spread in the distribution of expected utility might say: "This prospect has the same value to me as a prospect in which each were to have the same expectation—an expectation that is lower than 10 and greater than or equal to 4." If they are slightly sensitive to the spread in the distribution, then this expectation will be a number closer to 10. If they are more sensitive, then this expectation will be a number closer to 4.

What we need is a function with a single parameter α so that, as α moves from 0 to $+\infty$, the value of the distribution moves from 10 to 4. The following function does precisely this.³ For $i = 1, \dots, n$ persons, let x_i be the expectation of person i and, to respect anonymity, let us set the person-weights $w_i = 1/n$.

³ The function is in Atkinson 1970: 249-52 and is applied to prospects in Fleurbaey 2010: 658. This function is not rank-order sensitive. Alternatively, we can use a rank-order sensitive function as described in Donaldson and Weymark 1980: 74. See Bovens 2015a.

$$(1) \quad \begin{aligned} & \left(\sum_{i=1}^n w_i x_i^{(1-\alpha)} \right)^{1/(1-\alpha)} && \text{for } \alpha \neq 1 \\ & \exp\left(\sum_{i=1}^n w_i \ln x_i\right) && \text{for } \alpha = 1 \end{aligned}$$

This function yields the *equally distributed equivalent* of the distribution.

To illustrate, set α at the intermediate value of $\frac{1}{2}$. Then in our example,

$$(2) \quad \left(1/2\sqrt{16} + 1/2\sqrt{4} \right)^2 = 9$$

Since we apply the function to expectations here, it yields the *equally-distributed equivalent of the expectations* (EDEE) for a decision-maker with the inter-personal prospect distribution-sensitivity α_{EDEE} . So, we could say that the inter-personal prospective distribution-sensitive value of the prospect is not 10, but rather 9 for a decision-maker characterised by the sensitivity parameter $\alpha_{EDEE} = 1/2$.

Returning to our prospects *Risky* and *Certain* in Tables 1 and 2, the utilitarian value of the prospects is the average of the expectations, or, in other words, the equally distributed equivalent of the expectations with $\alpha_{EDEE} = 0$. Hence *Risky* ranks higher than *Certain* on the utilitarian calculus. However, the inter-personal prospect distribution-sensitive value with $\alpha_{EDEE} = 1/2$ of the prospects is:

$$(3) \quad V_{\alpha_{EDEE}=1/2}(\text{Risky}) = \left(1/2\sqrt{6.7} + 1/2\sqrt{3.4} \right)^2 = 4.91 \dots < 5 = V_{\alpha_{EDEE}=1/2}(\text{Certain})$$

Hence, a decision-maker with this type of sensitivity would choose *Certain* over *Risky*.

This is one distributional sensitivity. The decision-maker may bring other distributional sensitivities to the evaluation of the prospects. We will treat each of these distributional sensitivities in the same formal manner.

3.2.2 The Inter-State Distribution

Consider the inter-state distribution sensitivity. Suppose that there are two equiprobable states with average utilities 16 and 4. Then we could do the same exercise. The utilitarian is not sensitive to this distribution and the value of the prospect is just the expectation of the state averages, viz. 10. A decision-maker who is slightly sensitive to the inter-state distribution would consider the value of the prospect to be lower than but closer to 10, whereas a decision-maker who is very sensitive to the inter-state distribution would consider the value of the prospect to be closer to 4. The following function permits us to model inter-state sensitivities with x_j being state averages for states $j = 1, \dots, m$ and β ranging from 0 to $+\infty$.

$$(4) \quad \left(\sum_{j=1}^m p_j x_j^{(1-\beta)} \right)^{1/(1-\beta)} \quad \text{for } \beta \neq 1$$

$$\exp(\sum_{j=1}^m p_j \ln x_j) \quad \text{for } \beta = 1$$

This function yields the *certainty equivalent* of the distribution.

For β equals $\frac{1}{2}$, the value of the prospect with averages 16 and 4 in equiprobable states equals 9. Since we apply the function to state averages, it yields the *certainty equivalent of the averages* (CEA) for a person with inter-state sensitivity $\beta_{CEA} = \frac{1}{2}$.

Returning to our prospects *Risky* and *Certain*, the inter-state distribution-sensitive value of the prospects is:

$$(5) \quad V_{\beta_{CEA}=1/2}(Risky) = (.3\sqrt{11} + .7\sqrt{2.5})^2 = 4.42 \dots < 5 = V_{\beta_{CEA}=1/2}(Certain)$$

So, the decision-maker with this type of sensitivity would choose *Certain* over *Risky*.

3.2.3 The Intra-State Distribution

Suppose that the decision-maker is sensitive to the intra-state distribution. In this case, they will calculate the *equally-distributed equivalent* (EDE) of each state and set the value of the prospect at the expectation of the EDEs. Again, let us set $\alpha_{EDE} = 1/2$. In our example:

$$(6) \quad V_{\alpha_{EDE}=1/2}(Risky) = .3(1/2\sqrt{20} + 1/2\sqrt{2})^2 + .7(1/2\sqrt{1} + 1/2\sqrt{4})^2 = 4.17$$

which is smaller than $V_{\alpha_{EDE}=1/2}(Certain) = 5$.

3.2.4 The Intra-Personal Prospect Distribution

Finally, suppose that the decision-maker is sensitive to the intra-personal prospect distribution. Then they will calculate the *certainty-equivalents* (CE) for each person and set the value of the prospect at the average of the CEs. In our example with $\beta_{CE} = 1/2$:

$$(5) \quad V_{\beta_{CE}=1/2}(Certain) = 1/2(.3\sqrt{20} + .7\sqrt{1})^2 + 1/2(.3\sqrt{2} + .7\sqrt{4})^2 = 3.75$$

which is smaller than $V_{\beta_{CE}=1/2}(Risky) = 5$.

Table 3 provides an overview.

Distributional Sensitivity	Value of the Prospect	Parameter
Inter-Personal Prospect	Equally Distributed Equivalent of the Personal Prospect Expectations	α_{EDEE}
Inter-State	Certainty Equivalent of the State Averages	β_{CEA}
Intra-State	Expectation of the Equally Distributed Equivalents of the States	α_{EDE}
Intra-Personal Prospect	Average of the Certainty Equivalents of the Personal Prospects	β_{CE}

Table 3. Distributional Sensitivities, Prospect Values, and Parameters

4 Revisiting the Dilemmas

I will now revisit each of the dilemmas of Section 2. I will construct a stylised model for each dilemma and show how the distributional sensitivities underlie these dilemmas.

4.1 Sensitivity to the Inter-Personal Prospect Distribution and Restrictive Intervention

We model the case in which the people have been allocated to the groups that will and will not receive the more expensive drug. Suppose that without any intervention, there is a .70 survival rate. A cheaper intervention provided to all can increase the survival rate to .80. A more expensive intervention provided to half the population can increase the survival rate to $.90 + \varepsilon$. Let us first model the allocation without randomisation. For example, we might stipulate that the more expensive drug is provided to the urban population, but not to the equally large rural population.

<table border="1"> <thead> <tr> <th>EU</th> <th></th> <th></th> <th></th> <th></th> <th></th> </tr> </thead> <tbody> <tr> <td>.90</td> <td>0</td> <td>1</td> <td>1</td> <td>1</td> <td>...</td> </tr> <tr> <td>+ ε</td> <td></td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>...</td> <td>...</td> <td>...</td> <td>...</td> <td>...</td> <td>...</td> </tr> <tr> <td>.70</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>...</td> </tr> <tr> <td>...</td> <td>...</td> <td>...</td> <td>...</td> <td>...</td> <td>...</td> </tr> </tbody> </table>	EU						.90	0	1	1	1	...	+ ε					70	0	0	0	1	$>_{UTIL}$ $<_{InterPP}$	<table border="1"> <thead> <tr> <th>EU</th> <th></th> <th></th> <th></th> <th></th> </tr> </thead> <tbody> <tr> <td>.80</td> <td>0</td> <td>1</td> <td>1</td> <td>...</td> </tr> <tr> <td>...</td> <td>...</td> <td>...</td> <td>...</td> <td>...</td> </tr> <tr> <td>...</td> <td>...</td> <td>...</td> <td>...</td> <td>...</td> </tr> <tr> <td>...</td> <td>...</td> <td>...</td> <td>...</td> <td>...</td> </tr> <tr> <td>.80</td> <td>1</td> <td>1</td> <td>0</td> <td>...</td> </tr> </tbody> </table>	EU					.80	0	1	180	1	1	0	...
EU																																																																				
.90	0	1	1	1	...																																																															
+ ε																																																																				
...																																																															
.70	0	0	0	1	...																																																															
...																																																															
EU																																																																				
.80	0	1	1	...																																																																
...																																																																
...																																																																
...																																																																
.80	1	1	0	...																																																																
<i>Restrictive Intervention</i>		<i>Intervention for All</i>																																																																		

Table 4. *Restrictive Intervention* and *Intervention for All*

A utilitarian decision-maker prefers *Restrictive Intervention* to *Intervention for All*.

$$(8) \quad V_{UTIL}(RestrictiveIntervention) = .80 + \varepsilon/2 > .80 = V_{UTIL}(InterventionAll)$$

But now suppose that the decision-maker is sensitive to the inter-personal prospect distribution. For sufficiently large values of α_{EDDE} and sufficiently small values of ε , the value *Intervention for All* exceeds the value of *Restrictive Intervention* and hence they will prefer *Intervention for All*.

$$(9) \quad V_{\alpha_{EDDE} = x > 0}(\text{Restrictive Intervention}) < V_{\alpha_{EDDE} = x > 0}(\text{Intervention All})$$

In Ubel *et al.* 1996, the allocation is done by randomisation. The argument above still applies if we focus on the time-point after the randomisation has taken place. One might say that the medical ethicists and the prospective jurors, who favour *Intervention for All*, focus on the post-randomisation point when the persons in the prospect face unequal expectations in their personal prospects and are sensitive to the inter-personal prospect distribution. The decision theorists do not have this sensitivity and follow the utilitarian calculus.

The decision theorists could object that in case of randomisation, *Restrictive Intervention* is really the only reasonable solution. If we focus on the time point before the randomisation took place, each person's expectation from this programme is both (i) equal and (ii) greater than each person's expectation from the programme providing the cheaper drug to the whole population. Hence, the expensive drug should win out and sensitivity to the inter-personal prospect distribution cannot undo this.

The medical ethicists and prospective jurors could retort: Why should the preferred reference point be the pre-randomisation point? The decision theorist could respond that participants *would have* consented if they had been asked before the randomisation. But the fact of the matter is that they were not asked and why should *hypothetical* consent qualify as a justification? Someone who favours *Intervention for All* because they take the post-randomisation inequality to be morally relevant is not making a mistake or is not being irrational.

4.2 Sensitivity to the Inter-State Distribution and Diversification

Imagine a population in desperate need of a charity intervention. In the model, the outcome of a failed intervention is assigned the value 0 (e.g., for death) and the outcome of a successful outcome is assigned the value 1 (e.g., for full health). *GiveWell* recommends charities that are single cause and have a reasonable chance of success. The chance that the intervention is successful is, say, $2/3 + \varepsilon$. If the intervention is successful, then it is successful for all. *Oxfam* supports multiple causes. These interventions offer slightly less of a chance of success to each person—say, a chance of $2/3$. But they are sufficiently varied so that they won't all fail together. *Multiple-Cause* in Table 5 is represented as anti-correlated risk. (Here and in the tables below I represent prospects with anti-correlated risk rather than with independent risk. I do so for simplicity of presentation, because the contrast I wish to draw holds equally well for anti-correlated risk as for independent risk.)

				\succ_{UTIL} $\prec_{InterState}$				
EU	P(S ₁) = 1/3 - ε	P(S ₂) = 2/3 + ε			EU	P(S ₁) = 1/3	P(S ₂) = 1/3	P(S ₃) = 1/3
2/3 + ε	0	1	1	2/3	0	1	1	
2/3 + ε	0	1	1	2/3	1	0	1	
2/3 + ε	0	1	1	2/3	1	1	0	
SA	0	1	1	SA	2/3	2/3	2/3	
<i>Single-Cause</i>					<i>Multiple-Cause</i>			

Table 5. Single-Cause and Multiple-Cause

A utilitarian decision-maker prefers *Single-Cause* over *Multiple-Cause*:

$$(10) \quad V_{UTIL}(SingleCause) = 2/3 + \varepsilon > 2/3 = V_{UTIL}(MultipleCause)$$

The decision-maker who favours *Multiple-Cause* deviates from the utilitarian recommendation by adopting a sensitivity to the inter-state distribution. For a sufficiently large value of the parameter β_{CEA} and a sufficiently small ε ,

$$(11) \quad V_{\beta_{CEA} = x > 0}(SingleCause) < V_{\beta_{CEA} = x > 0}(MultipleCause)$$

They look at the state averages and are reluctant to support a single-cause charity that has a reasonable chance of being good for nothing. Deworming projects may have great value on the utilitarian calculus, but if the research on deworming does not bear out, as has been suggested in follow-up studies, then all the donations are money down the drain. The utilitarian value of single-cause charities may be greater, but the certainty equivalent of the state averages is greater for multiple-cause charities than for single-cause charities. In other words, they would rather play the short odds of Oxfam or MSF than the long odds of Deworming initiatives, even though the expected value of the latter is greater than of the former.

4.3 Sensitivity to the Intra-State Distribution and Unwelcome Risk Reductions

Let us construct a stylised model of the breast cancer screening case that uses different numbers than in Spiegelhalter's flow graph, but still captures the dilemma associated with screening for breast cancer in public health. Suppose that without screenings, the chances of survival are 2/3 and risks are independent. With screenings, there are no cancer fatalities. However, there are still costs to screenings: the mere imposition, the false positives and the unnecessary surgeries. We set the utility of screening at 2/3 - ε, so that each woman would opt for *No Screening* over *Screening*.

<table border="1"> <tr><th>EU</th><td></td><td></td><td></td></tr> <tr><td>2/3</td><td>0</td><td>1</td><td>1</td></tr> <tr><td>2/3</td><td>1</td><td>0</td><td>1</td></tr> <tr><td>2/3</td><td>1</td><td>1</td><td>0</td></tr> </table>				EU				2/3	0	1	1	2/3	1	0	1	2/3	1	1	0	$>_{UTIL}$ $<_{IntraState}$	<table border="1"> <tr><th>EU</th><td></td><td></td><td></td><td></td></tr> <tr><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td></tr> <tr><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td></tr> <tr><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td></tr> </table>				EU					2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ
				EU																																								
				2/3	0	1	1																																					
2/3	1	0	1																																									
2/3	1	1	0																																									
EU																																												
2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ																																								
2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ																																								
2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ																																								
<i>No Screening</i>	<i>Screening</i>																																											

Table 6. *No Screening* and *Screening*

A utilitarian decision-maker who is insensitive to the intra-state distribution prefers *No Screening* to *Screening*.

$$(12) \quad V_{UTIL}(NoScreening) > V_{UTIL}(Screening)$$

However, a decision-maker who is sensitive to the intra-state distribution dislikes the fact that, under *No Screening*, there is a wide spread in the utilities within each state. If the equally distributed equivalent of each state in *No Screening* drops below $2/3 - \epsilon$, then the value of the prospect *No Screening* will be lower than the value of the prospect *Screening*. This will indeed be the case for sufficiently large α_{EDE} and sufficiently small ϵ .

$$(13) \quad V_{\alpha_{EDE} = x > 0}(NoScreening) < V_{\alpha_{EDE} = x > 0}(Screening)$$

So, a decision-maker with a sensitivity for the intra-state distributions will overrule the preference of the women affected. To justify their decision, they will point to the fact that there is bound to be a large casualty rate within each state.⁴

4.4 Sensitivity to the Intra-Personal Prospect Distribution and Not on My Watch

Suppose that, given certain obstetric indications, an obstetrician must choose between vaginal birth, at the risk of the life of the mother and the foetus, or a C-section, which is safer but requires a longer recovery process. We model the stage at which a fully informed prospective mother would decline a C-section, but just barely; with a slight increase in the risk of harm to her or the foetus, the mother would opt for a C-section. For example, suppose that the mother has a 2/3 chance of the best outcome (say a successful vaginal birth) at utility 1 and a 1/3 chance of the worst outcome (say, death) at utility 0. If the mother chooses a C-section, she is certain to end up with utility $2/3 - \epsilon$.

Given that this concerns a pattern of decisions followed by obstetricians, we imagine that multiple women are in this situation.

⁴ A similar case and analysis is presented in Fleurbaey and Voorhoeve 2013.

<table border="1"> <tr><td>EU</td><td></td><td></td><td></td></tr> <tr><td>2/3</td><td>0</td><td>1</td><td>1</td></tr> <tr><td>2/3</td><td>1</td><td>0</td><td>1</td></tr> <tr><td>2/3</td><td>1</td><td>1</td><td>0</td></tr> </table>	EU				2/3	0	1	1	2/3	1	0	1	2/3	1	1	0	\succ_{UTIL} $\prec_{IntraPP}$	<table border="1"> <tr><td>EU</td><td></td><td></td><td></td></tr> <tr><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td></tr> <tr><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td></tr> <tr><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td><td>2/3 - ϵ</td></tr> </table>	EU				2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ
EU																																		
2/3	0	1	1																															
2/3	1	0	1																															
2/3	1	1	0																															
EU																																		
2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ																															
2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ																															
2/3 - ϵ	2/3 - ϵ	2/3 - ϵ	2/3 - ϵ																															
<i>Vaginal Birth</i>		<i>C-section</i>																																

Table 7. *Vaginal Birth* and *C-Section*

Since the expectation of vaginal birth is greater than of C-section, the affected prospective mothers will choose for a vaginal birth. If the prospective mother is herself an MD, then the obstetrician will simply exercise fiduciary care and execute her will.

$$(14) \quad V_{UTIL}(VaginalBirth) > V_{UTIL}(CSection)$$

If the prospective mother is not an MD, then the obstetrician will be deciding for her. Now suppose that the obstetrician is sensitive to the intra-personal prospect distribution. For a sufficiently large value of the parameter β_{CE} and a sufficiently small ϵ ,

$$(15) \quad V_{\beta_{CE} = x > 0}(VaginalBirth) < V_{\beta_{CE} = x > 0}(CSection)$$

When deciding for non-MD prospective mothers, they are deciding for others. Even though it may be reasonable for the prospective mother herself to take the risk of a vaginal birth, the obstetrician wishes to play it safe. In this case, playing it safe is playing it a bit safer than a reasonable and fully informed prospective mother would choose to play it herself. This can be modelled by assigning a value greater than 0 to the certainty equivalent parameter β_{CE} , measuring the decision-maker's sensitivity to the intra-personal prospect distribution.⁵

4.5 Summing Up

The first four dilemmas nicely map onto our four distributional sensitivities. The following table provides an overview.

⁵ Lara Buchak 2017: 22–24 argues that in making risky decisions for others when not knowing their risk attitude, we should operate with the most risk averse attitude within the range of what is reasonable, but if we know their risk attitude we should operate with their risk attitude. This could provide an alternative explanation. Since the obstetrician and the MD prospective mother have the same medical background, the obstetrician may be presumed to be more knowledgeable of the MD prospective mother's risk attitude. If the prospective mother is not an MD, then the obstetrician is less knowledgeable and decides for her with a more risk averse attitude.

Dilemma	Sensitivity for Distribution	Value of the Prospect
Restrictive Intervention	Inter-Personal Prospect	Equally Distributed Equivalent of the Personal-Prospect Expectations
Diversification	Inter-State	Certainty Equivalent of the State Averages
Unwelcome Risk Reductions	Intra-State	Expectation of the State Equally Distributed Equivalents
Not on My Watch	Intra-Personal Prospect	Average of the Personal-Prospect Certainty Equivalents

Table 8. Overview

In each dilemma, it is a particular sensitivity that provides a counter weight to the utilitarian recommendation. Someone who has this sensitivity will reasonably substitute the value of the prospect in the table for the utilitarian value of the prospect.

We now move on to the dilemma of *Saving Few at Greater Risk to Many* which is slightly more complex because it incorporates multiple distributional sensitivities.

5 Saving Few at Greater Risk to Many

5.1 Cases

Here are five cases of in which we must choose between saving the few or reducing the risk to the many.

- (i) We can either provide expensive anti-retrovirals to AIDS patients or increase efforts to prevent HIV transmission (e.g., through condom dispersal programmes).
- (ii) We can either treat those infected with Ebola or we can invest in prevention through vaccination projects.
- (iii) We can either invest millions on saving a miner who is trapped or direct the funds to improving mine safety so as to prevent future accidents.
- (iv) In the biblical Parable of the Lost Sheep (Luke 15: 3–7) we can go out and save one sheep while leaving the 99 sheep in the barn at the risk of a visit by wolves.⁶

⁶ Luke 15 is a response to the Pharisees who object that Jesus keeps the company of sinners and tax collectors. Jesus responds that there is more joy in saving people who have gone astray than in attending to the righteous. My interpretation may be a bit of a stretch as biblical

- (v) Diesel engines emit more NO_x and particulates. This imposes risk on present people who suffer from respiratory diseases. Fuel engines emit more CO₂ and impose risks on future people due to climate change and these risks will affect many. We can protect the present few by policies favourable to fuel engines, or we can reduce the risk to the future many by policies favourable to diesel engines.

There is a question as to whether the many are facing more or less correlated risk. The stories can be filled in in different ways. For some diseases that we fail to prevent, the risks to the many are independent. For other diseases, the risks are correlated; that is, there is the risk of an outbreak. Investment in mine safety may address more or less correlated risk. Depending on the hunting behaviour of the wolves, the risk to the sheep in the barn may be more or less correlated. In the case of climate change risks are clearly correlated.

All of these cases have the structure of a dilemma. In the case of independent risk, the pull comes from two directions. We want to save the most expected lives, yet it seems heartless not to save the few. In the case of correlated risk, there is an additional pull from a third direction, viz. we are wary of catastrophes in which things go wrong for the many.⁷

5.2 Modelling Saving Few at Greater Independent Risk to Many

Suppose that person n is sure to die unless we save them. If we forego saving them, then we can increase the survival chance for persons 1, ..., $n - 1$ from .95 to $.95 + \varepsilon$. Let us assume that these survival chances are independent. Under *Treatment*, resources are being allocated to person n who will now survive whereas otherwise they were sure to die. Under *Prevention*, we let person n die, but we reduce the risk to persons 1 through $n - 1$ who now face a better prospect. Suppose that $(n - 1)\varepsilon > 1$ so that the utilitarian decision-maker prefers *Prevention* to *Treatment*.

scholarship goes, but one could say that the cost of saving a few sinners is that one cannot give due attention to the righteous who are thereby put at increased risk of sinning.

⁷ The problem of saving few at the cost of a greater risk to many is discussed by Schelling 1968, Reibetanz 1998, Brock and Wikler 2009, Otsuka 2015, and Frick 2015. The Lost Sheep case is discussed in Bovens and Fleurbaey 2012 and Dieselgate is discussed in Bovens 2017.

<table border="1"> <tr><td>EU</td><td></td><td></td><td></td><td></td></tr> <tr><td>.95 + ε</td><td>0</td><td>1</td><td>1</td><td>...</td></tr> <tr><td>...</td><td>1</td><td>0</td><td>1</td><td>...</td></tr> <tr><td>...</td><td>...</td><td>...</td><td>...</td><td>...</td></tr> <tr><td>.95 + ε</td><td>1</td><td>1</td><td>0</td><td>...</td></tr> <tr><td>0</td><td>0</td><td>0</td><td>0</td><td>...</td></tr> </table>	EU					.95 + ε	0	1	1	1	0	195 + ε	1	1	0	...	0	0	0	0	...	$>_{UTIL}$ $<_{InterPP}$	<table border="1"> <tr><td>EU</td><td></td><td></td><td></td><td></td></tr> <tr><td>.95</td><td>0</td><td>1</td><td>1</td><td>...</td></tr> <tr><td></td><td>1</td><td>0</td><td>1</td><td>...</td></tr> <tr><td>...</td><td>...</td><td>...</td><td>...</td><td>...</td></tr> <tr><td>.95</td><td>1</td><td>1</td><td>0</td><td>...</td></tr> <tr><td>1</td><td>1</td><td>1</td><td>1</td><td>...</td></tr> </table>	EU					.95	0	1	1	...		1	0	195	1	1	0	...	1	1	1	1	...
EU																																																														
.95 + ε	0	1	1	...																																																										
...	1	0	1	...																																																										
...																																																										
.95 + ε	1	1	0	...																																																										
0	0	0	0	...																																																										
EU																																																														
.95	0	1	1	...																																																										
	1	0	1	...																																																										
...																																																										
.95	1	1	0	...																																																										
1	1	1	1	...																																																										
<i>Prevention</i>		<i>Treatment</i>																																																												

Table 9. *Prevention and Treatment*

However, if a decision-maker is sensitive to the distribution of the expectations (i.e. if they are sensitive to the inter-personal prospect distribution), then they will prefer *Treatment* to *Prevention*.

Formally, for the utilitarian decision-maker the value of the prospects is such that:

$$(16) \quad V_{UTIL}(Prevention) > V_{UTIL}(Treatment)$$

However, for sufficiently large α_{EDEE} and sufficiently small n and ε , the value of the prospects is such that:

$$(17) \quad V_{\alpha_{EDEE} = x > 0}(Treatment) > V_{\alpha_{EDEE} = x > 0}(Prevention)$$

Exactly the same reasoning applies here as in the case of *Restrictive Intervention*, discussed in Section 4.1.

5.3 Modelling Saving Few at Greater Correlated Risk to Many

Let us fill in the story of the Lost Sheep as a case of correlated risk. We call the action of saving the one sheep "*Saving One Sheep*" and attending to the sheep in the barn "*Protecting 99 Sheep*". *Saving One Sheep* saves the lost sheep on row 100, but at the expense of leaving the sheep on rows 1 through 99 and subjecting them to the risk of a visit from a pack of wolves. *Protecting 99 Sheep* gives up on the sheep on row 100 but makes sure that the sheep on rows 1 through 99 are properly protected.

<table border="1"> <tr><td>EU</td><td></td><td></td><td></td></tr> <tr><td>1</td><td>1</td><td>1</td><td>1</td></tr> <tr><td>...</td><td>...</td><td>...</td><td>...</td></tr> <tr><td>1</td><td>1</td><td>1</td><td>1</td></tr> <tr><td>0</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>SA</td><td>.99</td><td>...</td><td>.99</td></tr> </table>				EU				1	1	1	1	1	1	1	1	0	0	0	0	SA	.9999	$\succ_{InterState}$ $\prec_{InterPP}$	<table border="1"> <tr><td>EU</td><td></td><td></td><td></td><td></td></tr> <tr><td>.989...</td><td>1</td><td>...</td><td>1</td><td>0</td></tr> <tr><td>...</td><td>...</td><td>...</td><td>...</td><td>...</td></tr> <tr><td>.989...</td><td>1</td><td>...</td><td>1</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>...</td><td>1</td><td>1</td></tr> <tr><td>SA</td><td>1</td><td>...</td><td>1</td><td>.01</td></tr> </table>					EU					.989...	1	...	1	0989...	1	...	1	0	1	1	...	1	1	SA	1	...	1	.01
EU																																																															
1	1	1	1																																																												
...																																																												
1	1	1	1																																																												
0	0	0	0																																																												
SA	.9999																																																												
EU																																																															
.989...	1	...	1	0																																																											
...																																																											
.989...	1	...	1	0																																																											
1	1	...	1	1																																																											
SA	1	...	1	.01																																																											
<i>Protecting 99 Sheep</i>					<i>Saving One Sheep</i>																																																										

Table 10. *Protecting 99 Sheep* and *Saving One Sheep*

If the survival chances of the 99 sheep in the barn are $(99 - 1)/99 = .989\dots$, then the utilitarian decision-maker is indifferent:

$$(18) \quad V_{UTIL}(Protecting) = V_{UTIL}(Saving)$$

A decision-maker who is sensitive to the inter-personal prospect distribution objects to the fact that *Protecting 99 Sheep* imposes a wider spread of expectations on the sheep.

$$(19) \quad V_{\alpha_{EDEE} = x > 0}(Protecting) < V_{\alpha_{EDEE} = x > 0}(Saving)$$

A decision-maker who is sensitive to the inter-state distribution objects to the fact that *Saving One Sheep* leaves a wider spread between the average state utilities: Things may just go horribly wrong in *Saving One Sheep*.

$$(20) \quad V_{\beta_{CEA} = x > 0}(Protecting) > V_{\beta_{CEA} = x > 0}(Saving)$$

There is an intuitive story to be told here. On the one hand, we may be especially sensitive to the fact that some individuals have very low expectations. In this case we don't want to just give up on the sheep on row 100. Sensitivity to the inter-personal prospect distribution makes us more eager to save the few than a utilitarian decision-maker. On the other hand, we may particularly care about the possibility that things could just go horribly wrong for the collective. In this case, we should be prepared to give up on row 100 in order to avoid the possibility of a grand downfall for all. Sensitivity to the inter-state distribution makes us more eager to reduce the risk to the many than a utilitarian decision-maker.

So, in the case of correlated risk, the relevant sensitivities pull in opposite directions from the utilitarian ranking.

Each of these sensitivities has had its detractors over the years and debates in the literature about making risky decisions for others can be captured in terms of how much weight these sensitivities should receive in particular situations.

6.1 Against Inter-Personal Prospect Distribution Sensitivity

Here is a variant of a case presented by Reibetanz 1998. Suppose that 100 farmers will be ploughing a field. There is a bomb in the field and one is bound to blow themselves up in doing so. A uniquely qualified bomb expert can detonate the bomb safely but in so doing will incur pneumonia. We are comparing the following distributions. Utility 1 stands for full health, 0 stands for death, and .9 for incurring pneumonia.

				$>_{UTIL}$ $<_{InterPP}$					
1	1	1	1		.99	0	...	1	1
	1	1	1						
...
1	1	1	1		.99	1	...	1	0
.9	.9	.9	.9		1	1	...	1	1
Average Expectation: 100.9/101					Average Expectation: 100/101				
<i>Safe Detonation</i>					<i>Lethal Bomb</i>				

Table 11. *Safe Detonation* and *Lethal Bomb*

Safe Detonation, which has the greater utilitarian value, indeed seems to be the morally preferable option. However, if we are sufficiently sensitive to the inter-personal prospect-Distribution, that is for a sufficiently high α_{EDEE} , then the value of *Lethal Bomb* exceeds the value of *Safe Detonation*.

In this particular case, there are indeed various features that would make one hesitant about supporting the solution that is sensitive to the inter-personal prospect distribution:

- (i) The harm to the expert (pneumonia) is substantially smaller than the harm that is bound to befall one of the farmers (death).
- (ii) The risks are anti-correlated and the harm is sure to occur (unlike for independent risk).
- (iii) We are making a recommendation about what the expert should do, not about the policy that a decision-maker should implement affecting others.

Suppose that the expert is bound to lose a limb in carrying out their task, with the utility of this health state set at .50 (*Harmful Detonation*), or that they run a .50 chance of losing their lives (*Risky Detonation*). Suppose that the risks to the farmers are independent so that there is a small chance that nobody will be affected. Let us call this *Lethal Bomb**. And suppose that we are a decision-maker who needs to instruct the expert what to do. In this case, the

utilitarian calculus would still recommend *Harmful Detonation* and *Risky Detonation* over *Lethal Bomb**, since its average expectation equals $100.5/101 > 100/101$. But now, I submit, we would be much less comfortable following this recommendation.

6.2 Against Inter-State Distribution Sensitivity

Keeney 1980 studies prospects with binary utilities—say, either the people in the prospect will live or they will die. He shows that there is a conflict between the goal of risk equity, i.e. the goal of distributing risk equally between individuals, and the goal of catastrophe avoidance, i.e. minimizing the risk that a large number of individuals will be hit. Indeed, the best way to avoid a catastrophe is to concentrate the risk on a few individuals, but this sacrifice of a small group is anti-egalitarian.

Consider a case of independent risk. Compare the following two prospects with binary utilities. Suppose that one person is bound to die and the other is bound to live. Call this prospect *One Lives One Dies*. We can spread the risk so that each has an independent chance of dying of .50. Call this prospect *Coin Flip*. The certainty equivalent of the state averages of *One Lives One Dies* is .50. But for $\beta_{CEA} > 0$, the certainty equivalent of the state averages of *Coin Flip* is lower than .50. Hence, a decision-maker who is sensitive to the inter-state distribution will try to focus the risk on a single person at the cost of risk equity.

<table border="1"> <tr><th>EU</th><td></td><td></td></tr> <tr><td>0</td><td>0</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>1</td></tr> <tr><td>SA</td><td>1/2</td><td>1/2</td></tr> </table>	EU			0	0	0	1	1	1	SA	1/2	1/2	\sim_{UTIL} $\succ_{InterState}$	<table border="1"> <tr><th>EU</th><td></td><td></td><td></td><td></td></tr> <tr><td>1/2</td><td>0</td><td>1</td><td>0</td><td>1</td></tr> <tr><td>1/2</td><td>0</td><td>0</td><td>1</td><td>1</td></tr> <tr><td>SA</td><td>0</td><td>1/2</td><td>1/2</td><td>1</td></tr> </table>	EU					1/2	0	1	0	1	1/2	0	0	1	1	SA	0	1/2	1/2	1
EU																																		
0	0	0																																
1	1	1																																
SA	1/2	1/2																																
EU																																		
1/2	0	1	0	1																														
1/2	0	0	1	1																														
SA	0	1/2	1/2	1																														
<i>One Lives One Dies</i>		<i>Coin Flip</i>																																

Table 12. *One Lives One Dies* and *Coin Flip*

The same tension holds for correlated risk. Let us revisit the story of the Lost Sheep. The inter-state distribution sensitivity favours *Protecting 99 Sheep* over *Saving One Sheep*. *Protecting 99 Sheep* offloads the risk on one sheep, viz. the lost sheep that won't be saved. Risk equity is embodied in the inter-personal prospect distribution sensitivity which, as we saw, favours *Saving One Sheep*, since it spreads the risk from the lost sheep to the 99 sheep in the barn.

6.3 Against Intra-State Distribution Sensitivity

Adler 2012: 523–4 argues that a decision-maker who is sensitive to the intra-state distribution does not satisfy the Axiom of Weak Separability of Persons.

What does the axiom say? Take two prospects A and B with persons 1, ..., n. First, construct prospects A* and B* by adding a person who is completely unaffected: Whatever states

materialises, they receive a fixed utility u^* . Second, construct prospects $A^\#$ and $B^\#$ by adding a person who is also completely unaffected: Whatever state materialises, they receive a fixed utility $u^\#$. Certainly, the axiom says, whether we add one person with a certain prospect of u^* or another person with a certain prospect of $u^\#$ should not make a difference to the ranking of the prospects.

This axiom is violated if we calculate the value of prospects by means of $\alpha_{EDE} > 0$. Consider the prospects in Table 13. Set α_{EDE} at $\frac{1}{2}$. Start with the one-person prospects A and B; add one person with certain utility $u^* = 0$; and one person with certain utility $u^\# = .49$. Note how $A^* >_{\alpha_{EDE}=1/2} B^*$ but $B^\# >_{\alpha_{EDE}=1/2} A^\#$, in violation of the Axiom of Weak Separability of Persons.

<table border="1" style="margin: auto;"> <tr><td></td><td style="text-align: center;">S1</td><td style="text-align: center;">S2</td></tr> <tr><td style="text-align: center;">P1</td><td style="text-align: center;">0</td><td style="text-align: center;">1</td></tr> </table> <p style="text-align: center;">A</p>		S1	S2	P1	0	1	<table border="1" style="margin: auto;"> <tr><td></td><td style="text-align: center;">S1</td><td style="text-align: center;">S2</td></tr> <tr><td style="text-align: center;">P1</td><td style="text-align: center;">.49</td><td style="text-align: center;">.49</td></tr> </table> <p style="text-align: center;">B</p>		S1	S2	P1	.49	.49							
	S1	S2																		
P1	0	1																		
	S1	S2																		
P1	.49	.49																		
<table border="1" style="margin: auto;"> <tr><td></td><td style="text-align: center;">S1</td><td style="text-align: center;">S2</td></tr> <tr><td style="text-align: center;">P1</td><td style="text-align: center;">0</td><td style="text-align: center;">1</td></tr> <tr><td style="text-align: center;">P2</td><td style="text-align: center;">0</td><td style="text-align: center;">0</td></tr> </table> <p style="text-align: center;">$V_{\alpha_{EDE}=1/2}(A^*) = 1/8 = .1250$</p>		S1	S2	P1	0	1	P2	0	0	$>_{\alpha_{EDE}=1/2}$	<table border="1" style="margin: auto;"> <tr><td></td><td style="text-align: center;">S1</td><td style="text-align: center;">S2</td></tr> <tr><td style="text-align: center;">P1</td><td style="text-align: center;">.49</td><td style="text-align: center;">.49</td></tr> <tr><td style="text-align: center;">P2</td><td style="text-align: center;">0</td><td style="text-align: center;">0</td></tr> </table> <p style="text-align: center;">$V_{\alpha_{EDE}=1/2}(B^*) = .1225$</p>		S1	S2	P1	.49	.49	P2	0	0
	S1	S2																		
P1	0	1																		
P2	0	0																		
	S1	S2																		
P1	.49	.49																		
P2	0	0																		
<table border="1" style="margin: auto;"> <tr><td></td><td style="text-align: center;">S1</td><td style="text-align: center;">S2</td></tr> <tr><td style="text-align: center;">P1</td><td style="text-align: center;">0</td><td style="text-align: center;">1</td></tr> <tr><td style="text-align: center;">P2</td><td style="text-align: center;">.49</td><td style="text-align: center;">.49</td></tr> </table> <p style="text-align: center;">$V_{\alpha_{EDE}=1/2}(A^\#) = .4225$</p>		S1	S2	P1	0	1	P2	.49	.49	$<_{\alpha_{EDE}=1/2}$	<table border="1" style="margin: auto;"> <tr><td></td><td style="text-align: center;">S1</td><td style="text-align: center;">S2</td></tr> <tr><td style="text-align: center;">P1</td><td style="text-align: center;">.49</td><td style="text-align: center;">.49</td></tr> <tr><td style="text-align: center;">P2</td><td style="text-align: center;">.49</td><td style="text-align: center;">.49</td></tr> </table> <p style="text-align: center;">$V_{\alpha_{EDE}=1/2}(B^\#) = .49$</p>		S1	S2	P1	.49	.49	P2	.49	.49
	S1	S2																		
P1	0	1																		
P2	.49	.49																		
	S1	S2																		
P1	.49	.49																		
P2	.49	.49																		

Table 13. Adler's Objection

How problematic is this? The axiom *seems* attractive because we confuse being unaffected with being uninvolved in the project. Indeed, it would be problematic if an uninvolved person could just be tagged onto the prospect and the ordering would flip depending on whether we tag on an uninvolved homeless person or an uninvolved millionaire. Uninvolved people should not be included. However, the fact that one is unaffected by a given policy decision does not mean that one is uninvolved. Involvement is determined by one's place in the social world. A choice between remuneration policies in the company where I am employed may not affect me in that it does not make a difference to *my* income. However, I am involved on

grounds of being an employee of the company and the choice between prospects should feature me.

Whether an unaffected person involved in a prospect is richer or poorer does make a difference to the intra-state distribution. We posit a decision-maker who is sensitive to the intra-state distribution. In the choice between A and B this sensitivity does not enter in, since these are one-person prospects. The decision-maker prefers A to B, since it offers a higher expectation. Adding a poor person to the prospect does not make B* more attractive to the decision-maker and they continue to prefer A*. Adding a person with the same utility as the person in B introduces equality into B#, which makes it into a more attractive option than A# to our decision-maker.⁸

6.4 Against the Intra-Personal Prospect Distribution

Otsuka and Voorhoeve 2018, elaborating on an argument proposed earlier in Otsuka and Voorhoeve 2009, construct the following case, modelled on Nagel 1979. There is a fifty-fifty chance that your child will turn out to be disabled or not. If they are disabled, then they will be slightly better off in town than in the country. If they are healthy, then they will be much better off in the country than in town. Should you move to town or to the country? We set the utilities as follows:

Disabled living in the Country	Disabled living in Town	Healthy living in Town	Healthy living in the Country
0	.10	.80	1

Table 14. Utilities in Otsuka and Voorhoeve’s Example

A single child, if idealized to be rational and well-informed, would choose to move to the country since it yields greater expected welfare. The question is whether a decision-maker choosing for the child could overrule this choice. Otsuka and Voorhoeve believe that a decision-maker should opt for the prospect that maximises the child’s expected welfare. If there are two children in the prospect facing fully correlated risks, then the same holds. But if they are facing fully anti-correlated risk, then we should choose for *Town*. See the rankings for O&V in Table 15.

Otsuka and Voorhoeve are sensitive to the intra-state distribution. They favour *Town* in the case of anti-correlated risk because the health inequalities within each state are less pronounced than in *Country*. However, they are not sensitive to the intra-personal prospect distribution. We can only overrule the choice of the persons in the prospects on grounds of intra-state inequalities.

Parfit 2012 disagrees. He argues that the parents should choose *Town* in each case. His argument is prioritarian: It is better to procure the smaller benefit of $.1 - 0 = .1$ by being located in town for a child that turns out to be disabled than the larger benefit of $1 - .8 = .2$

⁸ See also Voorhoeve and Fleurbaey 2016: 943–4.

by being located in the country for a child that turns out to be healthy. He objects to Otsuka and Voorhoeve that whether risks are correlated or anti-correlated cannot possibly make a difference to the moral ranking over the prospects. The fact that my interests matter more when I am disabled is morally significant as such and not because I am poorly off relative to someone else being better off.

Otsuka and Voorhoeve disagree. In the case of anti-correlation, one sibling will turn out healthy and one sibling will turn out disabled. Then by moving to the country, we know that a situation will ensue in which the disabled sibling can make the following forceful complaint: “Dear Parents, I have a reasonable complaint about your moving us to the country. It was a benefit to my sibling, but a harm to me, and even though the harm to me is smaller than the benefit to my sibling, you should have given priority to whoever would turn out to be worse off.” The same cannot be said in the case of correlated risk.

That is true, but something very similar can be said by the siblings if they both end up being disabled. They may complain: “Dear Parents, we have a reasonable complaint about your moving us to the country. It is a harm to us. Granted, the harm is smaller than the benefit that it would have offered us had we been healthy. However, you should have given more weight in your decision-making to the harm in case we would turn out to be disabled than the benefit in case we would turn out to be healthy.” Why would this counterfactual comparison with counterpart siblings in the case of correlated risk provide any less of a forceful complaint than the actual comparison between siblings in the case of anti-correlated risk?

I argued earlier that it is perfectly acceptable for a decision-maker who is sensitive to the intra-personal prospect distribution to add some risk aversion to the risk attitudes of the persons in the prospect. The parents of our siblings may argue that the risk of moving to the country is a risk that they cannot take for others—they can’t have the child end up disabled in the country *on their watch*. Sensitivity to the intra-personal prospect distribution can justify a choice *or Town over Country* in each of the three cases of Table 15. Hence, in this case, sensitivity to the intra-personal prospect distribution could support the same rankings as Parfit’s prioritarian argument.

<table border="1"> <tr><td>Disabled</td><td>Healthy</td></tr> <tr><td>.1</td><td>.8</td></tr> </table>	Disabled	Healthy	.1	.8	$\prec_{O\&V}$ \succ_{Parfit} $\succ_{IntraPP}$	<table border="1"> <tr><td>Disabled</td><td>Healthy</td></tr> <tr><td>0</td><td>1</td></tr> </table>	Disabled	Healthy	0	1
Disabled	Healthy									
.1	.8									
Disabled	Healthy									
0	1									
<i>Town</i>		<i>Country</i>								

<table border="1"> <tr><td>Disabled</td><td>Healthy</td></tr> <tr><td>.1</td><td>.8</td></tr> <tr><td>.1</td><td>.8</td></tr> </table>	Disabled	Healthy	.1	.8	.1	.8	$\prec_{O\&V}$ \succ_{Parfit} $\succ_{IntraPP}$	<table border="1"> <tr><td>Disabled</td><td>Healthy</td></tr> <tr><td>0</td><td>1</td></tr> <tr><td>0</td><td>1</td></tr> </table>	Disabled	Healthy	0	1	0	1
Disabled	Healthy													
.1	.8													
.1	.8													
Disabled	Healthy													
0	1													
0	1													
<i>Town</i>		<i>Country</i>												

<table border="1"> <tr><td>Disabled</td><td>Healthy</td></tr> <tr><td>.1</td><td>.8</td></tr> <tr><td>.8</td><td>.1</td></tr> </table>	Disabled	Healthy	.1	.8	.8	.1	$\succ_{O\&V}$ \succ_{Parfit} $\succ_{IntraPP}$	<table border="1"> <tr><td>Disabled</td><td>Healthy</td></tr> <tr><td>0</td><td>1</td></tr> <tr><td>1</td><td>0</td></tr> </table>	Disabled	Healthy	0	1	1	0
Disabled	Healthy													
.1	.8													
.8	.1													
Disabled	Healthy													
0	1													
1	0													
<i>Town</i>		<i>Country</i>												

Table 15. Otsuka and Voorhoeve's Objection to Prioritarianism

7. Discussion

We have discussed four distributional sensitivities that may, in certain decision situations, upset the utilitarian ordering. Multiple sensitivities can enter in at the same time and either reinforce each other or pull in opposite directions. In Bovens 2015a and 2015b, I proposed a model that integrates these sensitivities.

It is a mistake to think that there is a single model that will tell us once and for all how we should resolve dilemmas of risky decision-making for others. The information that is contained in prospects is too impoverished to determine a unique ranking. To determine a moral ranking over risky prospects affecting others we need to have access to much more detail than matrices of utilities with persons in the rows, states in columns and probability distributions over states.

The situation is comparable to Sen's 1970, 1976 and Gibbard's 1974 discussion of the Libertarian Paradox. In Sen's version of the paradox, Prude and Lewd find a copy of *Lady Chatterley's Lover* in the book store. Lewd thinks it would be a shame for the book to remain unread. She would like to read it herself. But she prefers that Prude read it, because it would do him some good. Prude would rather have it that nobody read the book. But if someone is

going to read it, then it better be him, because heaven knows what it might do to Lewd. Table 16 contains the preference orderings.

Prude	Lewd
O	P
P	L
L	O

Table 16. Sen’s Libertarian Paradox

Who, if anyone, should walk out with the book? Again, we have a dilemma. The Principle of Minimal Libertarianism says that one should be decisive about at least some options in one’s private sphere, such as, whether to read or not to read a particular book. Prude and Lewd should each be decisive as to whether they will or will not read the book. Hence, $L > O$ and $O > P$ and so, by Minimal Libertarianism, we should hand the book to Lewd. But note that both Lewd and Prude prefer $P > L$. So, by the Pareto Principle, we should hand the book to Prude. What should we do?

Gibbard presents a version of Gilbert and Sullivan’s *Trial by Jury* that has a similar structure. It is not quite respectful of modern sensitivities, but as a comedy and a critique of mores, it still passes muster. Edwin would rather not marry Angelina, but if she is going to marry the Judge, then he would rather marry her than see her marry the Judge. Angelina would rather marry Edwin than the Judge, but she prefers marrying the Judge to not marrying anyone.

Edwin	Angelina
O	E
E	J
J	O

Table 17. Gibbard’s version of the Libertarian Paradox

Who, if anyone, should get married to whom? Minimal Libertarianism stipulates the following. If the Judge is a willing party, then it is in Angelina’s private sphere whether to marry him or not. If Angelina is a willing party, then it is in Edwin’s private sphere whether to marry her or not. So, $J > O$ and $O > E$. By Minimal Libertarianism, we recommend that the Judge marry Angelina. However, both Edwin and Angelina prefer to be married to each other than for Angelina to be married to the Judge. Hence, by Pareto, $E > J$. What should be our recommendation?

In Sen’s version, there is something to be said for discounting nosy preferences and respecting each person’s desire to read or not read. At least, this is how Sen sees the case. Hence, Minimal Libertarianism wins and the book goes to Lewd. In Gibbard’s version, there is something to be said for letting the parties express their preferences. Angelina can threaten Edwin to marry the Judge, and Edwin will come to realise that he can only avoid this outcome by marrying Angelina. Hence, Pareto wins and Angelina and Edwin will tie the knot. At least, this is how Gibbard sees the case.

In Sen's response to Gibbard, he concedes that there is no single solution to the Libertarian Paradox. To resolve the dilemma in a morally sensitive matter, we need much more information than the information that is contained in the preference profile. The preference profile is too informationally impoverished to determine whether Pareto or Minimal Libertarianism should have the overhand.

Similarly, prospects alone are too informationally impoverished to determine a unique ordering. Different distributional sensitivities are fitting in different contexts. And just as in Sen's and Gibbard's Libertarian Paradoxes, different dilemmas may be represented by the same prospects and yet demand very different solutions.

Here is an example of very different types of decisions that can be modelled by the same pair of prospects. Let us revisit the case of screening for breast cancer which we modelled in Table 6. Public health dilemmas over alcohol policy can be modelled by the same pair of prospects. There are various things one can do to reduce alcohol consumption—increased taxes, restrictions on deep discounting (that is, a sale price below VAT and excise taxes), minimal unit pricing, limiting sales times etc. Many people object to restrictive alcohol policies. However, when we consider the number of alcohol-related deaths prevented in the population at large by restrictive policies, one might think that this annoyance is a small price to pay for the benefits gained. The stylised prospects for *No Screening* and *Screening* could also serve as models for respectively a liberal alcohol policy and a restrictive alcohol policy.

A decision-maker with a sensitivity for the intra-state distribution might defend a policy of intensive screening for breast cancer, overruling the preferences of a majority of the women involved. We know that it is highly likely that many more women would die with less intensive screening and from a public health standpoint saving these women through intensive screening is worth the annoyance and unnecessary treatment. I have some sympathy for this position. However, in the case of alcohol policy, I am less sympathetic to restrictive alcohol policies. Certainly, we can save many lives by introducing more restrictive alcohol policies, and yet, I am less inclined to overrule the preferences of a population that resists such moves. In England, suggestions of more restrictive alcohol policies are met with the defiant response: "We do not want to become Sweden."

So even though the pair of prospects that model both public health dilemmas are the same, I am more inclined to let the Pareto condition be overruled in the case of breast cancer screening than in the case of alcohol policies. What makes for the difference? In the case of breast cancer screening who will and will not be affected by breast cancer is down to brute luck. But in the case of alcohol policies, who will and will not be affected by alcohol-related morbidity is to a certain extent down to the choices that people make. For this reason, I am less willing to overrule the preference of many currently responsible drinkers who want to continue enjoying their pint English-style. This distinction is not reflected in prospects: All we have is a probability distribution over better or worse consequences that may come about under different screening regimes or different alcohol policies. Just as in the Sen-Gibbard debate, there is information that is extraneous to the formal model, viz. whether people have more or less control over health outcomes, that determines how much weight conflicting reasons merit in resolving the dilemma.

What is the good of modelling, if not to determine a single decision algorithm ideally backed up by an axiomatic justification? What progress did we make by modelling these multiple sensitivities? Making risky decisions for others in a morally responsible matter is not an algorithmic exercise and the formal machinery is not meant to make it such. However, formal machinery helps us understand what forces are at work when we find ourselves in a quandary and what is driving people when they disagree in the forum. This is an important contribution to moral decision-making. It is a mistake to expect more definite answers from modelling.⁹

References

Adler, Matthew. 2012. *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. Oxford: Oxford University Press.

Aiken, Alexander M., Calum Davey, James R. Hargreaves, and Richard J. Hayes. 2015. "Re-analysis of health and educational impacts of a school-based deworming programme in western Kenya: a pure replication." *International Journal of Epidemiology* 44: 1572–1580.

Atkinson, Anthony B. 1970. On the Measurement of Inequality. *Journal of Economic Theory* 2: 244–263.

Bovens, Luc and Marc Fleurbaey. 2012. "Evaluating Life or Death Prospects." *Economics and Philosophy* 28: 217–49.

Bovens, Luc. 2015a. "Concerns for the poorly off in ordering risky prospects." *Economics and Philosophy* 31: 397–492.

Bovens, Luc. 2015b. "Evaluating Risky Prospects: the Distribution View." *Analysis* 75: 243–53.

Bovens, Luc. 2017. "The Ethics of Dieselgate." *Midwest Studies in Philosophy* 40:262-83.

Brock, Dan W. and Daniel Wikler. 2009. "Ethical Challenges in Long-Term Funding For HIV/AIDS." *Health Affairs* 28: 1666–76.

Buchak, Lara. 2017. "Taking Risks Behind the Veil of Ignorance." *Ethics* 127: 1–35.

Davey, Calum, Alexander M Aiken, Richard J. Hayes, and James R. Hargreaves. 2015. "Re-analysis of health and educational impacts of a school-based deworming programme in western Kenya: a statistical replication of a cluster quasi-randomized stepped-wedge trial." *International Journal of Epidemiology*, 44: 1581–1592.

Donaldson, David and John A. Weymark. 1980. "A Single-Parameter Generalization of the Gini Indices of Inequality." *Journal of Economic Theory* 22: 67–86.

⁹ I am grateful for comments on earlier drafts by David Kinney, Ruth Tomlin, Alex Voorhoeve and Mark White.

Fleurbaey, Marc. 2010. "Assessing risky social situations." *Journal of Political Economy* 118: 649–680.

Fleurbaey, Marc and Alex Voorhoeve. 2013. "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, Measurement, and Ethics*. (Oxford: Oxford University Press), pp. 113–28.

Frick, Johann. 2015a. "Contractualism and Social Risk." *Philosophy and Public Affairs* 43: 175–223.

Frick, Johann. 2015b. "Treatment versus Prevention in the Fight Against HIV/ AIDS and the Problem of Identified versus Statistical Lives." In *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. Ed. Glenn Cohen, Norman Daniels, & Nir Eyal. New York: Oxford University Press.

Gauthier, David. 1963. *Practical Reasoning: The Structure and Foundations of Prudential and Moral Arguments and Their Exemplification in Discourse*. Oxford: Clarendon Press.

Gibbard, Alan. 1974. "A Pareto-Consistent Libertarian Claim." *Journal of Economic Theory* 7: 388–410.

GiveWell. URL: <http://www.givewell.org>. Accessed on 2 February 2018.

Greaves, Hilary. 2017. "A reconsideration of the Harsanyi-Sen-Weymark debate on utilitarianism." *Utilitas* 29: 175–213.

Harsanyi, John. 1955. "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility." *Journal of Political Economy* 63: 309–21.

Johnson, Erin M., and M. Marit Rehavi. 2016. "Physicians Treating Physicians: Information and Incentives in Childbirth." *American Economic Journal: Economic Policy* 8: 115–41.

Keeney, Ralph L. 1980. "Equity and Public Risk." *Operations Research* 28: 527–34.

Miguel, Edward and Michael Kremer. 2003. "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities." *Econometrica* 71: 159–217.

Nagel, Thomas. 1979. "Equality." In Thomas Nagel *Mortal Questions*. (Cambridge: Cambridge University Press), pp. 106–27.

Moreau, Sophia Reibetanz. 1998. "Contractualism and Aggregation." *Ethics* 108: 296–31.

Otsuka, Michael and Alex Voorhoeve. 2009. "Why It Matters that Some Are Worse Off Than Others: An Argument Against the Priority View." *Philosophy & Public Affairs* 37: 171–99.

Otsuka, Michael. 2015. "Risking Life and Limb. How to Discount Harms by Their Improbability." In Glenn Cohen, Norman Daniels, & Nir Eyal (eds.) *Identified Versus Statistical Lives: An Interdisciplinary Perspective* (Oxford: Oxford University Press), pp. 77–93.

Otsuka, Michael and Alex Voorhoeve. 2018. "Equality and Priority." In Serena Olsaretti (ed.) *Oxford Handbook of Distributive Justice*. (Oxford: Oxford University Press).

Parfit, Derek. 2012. "Another Defence of the Priority View." *Utilitas* 24: 399–440.

Rawls, John. 1974–75. "The Independence of Moral Theory." *Proceedings and Addresses of the American Philosophical Association*. 47. Reprinted in Rawls, John. 1999. *Collected Papers*. Samuel Freeman (ed.) (Boston MA: Harvard University Press) pp. 286–302.

Reibetanz, Sophia. 1998. "Contractualism and Aggregation." *Ethics* 108: 296–331.

Schelling, Thomas. 1968. "The Life You Save May Be Your Own." In Samuel B. Chase (Ed.) *Problems in Public Expenditure Analysis. Studies of Government Finance*. (Washington: The Brookings Institution.) Reprinted in Thomas Schelling. 1984. *Choice and Consequence*. (Boston MA: Harvard University Press) pp. 113–46.

Sean 2017 "Deworming might have huge impact, but might have close to zero impact." GiveWell Blog. URL: <https://blog.givewell.org/2016/07/26/deworming-might-huge-impact-might-close-zero-impact/>. Accessed on 2 February 2018.

Sen, Amartya. 1970. "The Impossibility of a Paretian Liberal." *Journal of Political Economy* 78: 152–57.

Sen, Amartya. 1973. *On Economic Inequality*. (Oxford: Clarendon Press).

Sen, Amartya. 1976. Liberty, Unanimity and Rights. *Economica* 43: 217–45.

Spiegelhalter, David. 2015. "A Visualisation of the Information in NHS Breast Cancer Screening Leaflet." URL: <https://understandinguncertainty.org/visualisation-information-nhs-breast-cancer-screening-leaflet>. Accessed on 2 February 2018.

Taylor-Robinson David C., Nicola Maayan, Karla Soares-Weiser, Sarah Donegan, and Paul Garner. 2015. Deworming drugs for soil-transmitted intestinal worms in children: effects on nutritional indicators, haemoglobin, and school performance. *Cochrane Database of Systematic Reviews*. DOI: 10.1002/14651858.CD000371.pub6.

Ubel, Peter A., Michael L. DeKay, Jonathan Baron, and David A. Asch. 1996. "Cost-Effectiveness Analysis in a Setting of Budget Constraints." *The New England Journal of Medicine* 334: 1174–77.

Voorhoeve, Alex and Marc Fleurbaey. 2016. "Priority or Equality for Possible People." *Ethics* 126: 929–54.

Weymark, John A. 1991. "A Reconsideration of the Harsanyi-Sen Debate on Utilitarianism." In Jon Elster and John Roemer (eds.) *Interpersonal Comparisons of Well-Being*. (Cambridge: Cambridge University Press) pp. 255–320.

Williams, Bernard. 1973. "A Critique of Utilitarianism." In *Utilitarianism: For and Against*. (Cambridge: Cambridge University Press).