

Alex Voorhoeve, draft a25, March 30, 2021. Forthcoming in Conrad Heilmann and Julian Reiss (eds.) *The Routledge Handbook of Philosophy of Economics*.

Policy Evaluation under Severe Uncertainty: A Cautious, Egalitarian Approach

Word count (excl. abstract, references and biographical note): 6,754

Abstract: In some severely uncertain situations, exemplified by climate change and novel pandemics, policymakers lack a reasoned basis for assigning probabilities to the possible outcomes of the policies they must choose between. I outline and defend an uncertainty averse, egalitarian approach to policy evaluation in these contexts. The upshot is a theory of distributive justice which offers especially strong reasons to guard against individual and collective misfortune.

Introduction

When policymakers evaluate a policy, they are typically unsure what would result from choosing it. In welfare economics, such lack of knowledge is commonly dealt with by: (i) assigning precise probabilities to the possible outcomes of every policy under evaluation; (ii) assigning a value to each of these possible outcomes, for example by using a social welfare function which evaluates the distribution of well-being in each possible outcome; and finally (iii) recommending the policy with the highest expected value (the probability-weighted sum of the values of its possible outcomes) (Fleurbaey 2010; Adler 2019). Here, I will follow common practice decision theory and call situations in which a decisionmaker is in a position to make such expected value calculations “decision problems under risk.”

Sometimes, however, those who decide on a policy are not able to compute these expected values for a policy because they are unable to non-arbitrarily assign precise probabilities to every one of its possible outcomes. Here, I shall follow Knight (1921) and Keynes (1937) and refer to such situations as calculations as “uncertain.” (Following Ellsberg [1961], such situations are also commonly referred to as “ambiguous”.) Throughout, I use both “risk” and “uncertainty” in their subjective senses—as pertaining to the beliefs about the chances of possible outcomes of their decisions that a rational decision-maker can form on the basis of their priors and the evidence available to them. (For further discussion of the contrast between risk and uncertainty, see Stefansson’s chapter in this volume.)

In welfare economics and political philosophy, there has been far more discussion of how to do expected value evaluation of policies than of how to make public decisions under uncertainty. In this chapter, I take a step towards remedying this comparative neglect. This project is worth pursuing because severely uncertain situations are common and important. One example is climate change (Heal and Millner 2018). In judging climate policies, policymakers face both climate-scientific uncertainty—about how the climate system works and would react to various emissions scenarios—and socio-economic uncertainty—about how individuals and societies will respond to changes in climate. Climate-scientific uncertainty arises because our main source of predictions about what might happen in various emissions scenarios are climate models. These are sensitive to changes in specified initial conditions which are only imperfectly known. They are also highly sensitive to the choice of functional forms by which key relationships are represented, and of key parameter values, both of which are also imperfectly known (Frigg et al. 2014). There is also scientific disagreement about some key causal mechanisms, how important they are, and how they

interact with widely accepted mechanisms. Because of the diversity of these sources of uncertainty, some of which cannot be captured by probability distributions over different potential initial values, functional forms, parameter values, and causal mechanisms, and because these different sources interact in the context of extremely sensitive models, it is commonly thought that it is not possible to non-arbitrarily capture current climate scientific knowledge by assigning a precise probability to key propositions, such as that in a “medium emissions scenario,” the Earth will warm by more than 2.0 degrees centigrade (Dietz 2014; Heal and Millner 2018). For this reason, the most authoritative report available, by the Intergovernmental Panel on Climate Change (IPCC) assigns only ranges of probabilities to such propositions. For example, it reports that in one medium emissions scenario, “warming is *likely* to exceed 2.0 degrees centigrade,” where “likely” means: “has a probability of between 66% and 100%” (IPCC 2014, p. 10).

Social-scientific uncertainty about the impact on societies and people’s lives of various degrees of warming is arguably greater still (Heal and Millner 2018). For there is very little evidence about how changes in climate might affect such things as political stability, migration flows, or economic growth. Consequently, experts regard socio-economic impact assessments as highly speculative (Dietz 2014; Heal and Millner 2018). It follows that it would represent a leap beyond the information available to assign precise probabilities to outcomes of interest, such as “there is 3.0 degrees warming, a permanent loss of economic output of 5% of GDP, and there are large forced migration flows.”

A second example of decision making under severe uncertainty is presented by novel pandemics. Again, the sources of uncertainty are multifarious. In the early months of the COVID-19 pandemic, for example, there was lack of information and there were grave

differences of expert opinion about such key variables as the transmissibility of the virus that causes COVID-19, its infection fatality rate, and the potential effectiveness of novel treatments and vaccines. There was also a notable lack of information and consensus about how various social and movement measures (e.g., mask wearing and “lockdowns”) might be expected to impact the virus’ spread and, more broadly, health and other components of well-being. A significant source of information for policy decisions in the early months were non-peer-reviewed models of disease spread. These models are non-linear, and their outcomes are highly dependent on assumed initial conditions and parameter values. Because of this sensitivity, small differences in these assumptions readily generated outcomes of interest (e.g., “deaths from disease with a given policy”) that differed by several orders of magnitude (Avery et al. 2020, pp. 10-11, and p. 20). Moreover, these assumptions were highly uncertain: initial conditions were unknown, functional forms disputed, and parameter values were often *ad hoc*. These diverse sources of uncertainty meant that there was no clear basis to assign precise probabilities to even one aspect of interest in evaluating policies, namely their impact on the spread of disease and associated deaths. Indeed, among five models of possible deaths in the UK and USA that achieved prominence in policy discussions early in the pandemic, only one provided probabilities for its estimates, and those were dubious (Avery et al. 2020, p. 26). The quality of information on possible wider health, social and economic impacts of social and movement measures such as lockdowns was, at the time, equally poor, in part because of the lack of precedent for such measures in contemporary economies. Assigning precise probabilities to the possible outcomes of key policies would therefore have represented a leap beyond the available evidence.

In this chapter, I outline an approach to policy evaluation for such uncertain situations. There is political-philosophical work to be done in using the tools of decision theory for this

purpose. For the bulk of the literature has been devoted to the question how, under uncertainty, people actually do and rationally may make decisions *on their own behalf*. There is less work on how to make *public* decisions in the absence of precise probabilities.¹

I proceed as follows. In Section 1, I outline a pluralist egalitarian theory of distributive justice for situations of risk that I will take as my point of departure. In Section 2, I make the case for the permissibility of using a cautious decision criterion under uncertainty. In Section 3, I explore some implications of incorporating this form of caution into the outlined egalitarian view. I show that caution strengthens one element of egalitarian solidarity by reinforcing its concern for those who may end up worse off than others. But I also show that it may counteract another element of such solidarity, namely the tendency to ensure that everyone “sinks or swims” jointly. In Section 4, I conclude.

1. Egalitarianism under Risk

On the pluralist egalitarian view that I shall draw on here, people’s interests should be considered from two perspectives. The first is in terms of the value of their prospects. A person’s prospects are important because they capture the extent to which a policymaker’s actions promote this person’s interests as they are rationally viewed with the information at hand when deciding on a policy. The second is in terms of each person’s final well-being. This is relevant because this represents the interests that a policymaker should aim to see advanced equally, if they were fully informed about how each person would end up faring.

On the proposed egalitarian view, it is unfair when people’s interests in having good prospects and in faring well are advanced unequally. Besides limiting these two forms of

inequality, the view in question is also concerned with promoting people's interests, both in prospect and in terms of their final well-being. In sum, on this view, a policy-maker should adopt the following aims: (i) reducing inequalities in the value of people's prospects; (ii) reducing inequality in final well-being; (iii) improving people's prospects; and (iv) improving final well-being.²

Throughout, I assume an interpersonally comparable, cardinal measure of well-being derived from idealized preferences under risk. On this measure, a first policy yields higher expected well-being for a person than a second policy just in case it would be preferred for this person's sake after rational deliberation with relevant knowledge while taking into account only this person's self-interest. A first policy yields the same expected well-being as a second policy for a person if and only if such a deliberator would be indifferent between the two policies on this person's behalf. (For a defense of this measure, which is common in parts of welfare economics see Adler 2019, Appendix D.) I shall also assume that even though individuals' actual preferences risk may diverge from this idealization because of reasoning errors and biases, individuals accept the idea that their good should be measured by the preferences they would have after deliberation that corrected these errors and biases, so that in measuring their well-being by these idealized preferences, we are aligning ourselves with their judgments (Arneson 1990). Finally, I assume that in situations of risk it is permissible for a policymaker to maximize expected moral value. In sum, I adopt an orthodox approach under risk. This allows me to focus on the departure I make from the orthodoxy in cases of severe uncertainty.

To illustrate the outlined egalitarian view under risk, imagine that you are a resource allocation manager in a government-run health system. Two twenty year-old citizens, Ayan

and Bashir, face a debilitating illness which, if untreated, will leave them unable to walk and so limited in dexterity that they will require the help of another person for most tasks.

Consequently, they will have a lifetime well-being of 30 (a merely tolerable quality of life). If fully cured, each would have a lifetime well-being of 80 (a very good quality of life). You do not have enough resources to fully cure them both for sure. Instead, you can allocate resources towards one of the alternatives outlined in the top section of Table 1. To link up with Ellsberg's (1961) paradigmatic contrasting presentations of risky and uncertain alternatives, risk will be represented by a random draw from an urn which is known to contain only 50 red balls and 50 black balls. The numbers in brackets in the table are the probabilities associated with each draw.

Table 1. Final Well-being for All Alternatives

Risky Alternatives	Draw from a risky urn	
	<i>Red (0.5)</i>	<i>Black (0.5)</i>
<i>Inequality under Certainty</i>		
Ayan	80	80
Bashir	30	30
<i>Equal Risk, Unequal Final Well-being</i>		
Ayan	80	30
Bashir	30	80
<i>Equality under Certainty</i>		
Ayan	55	55
Bashir	55	55
<i>Equality under Risk</i>		
Ayan	80	30
Bashir	80	30
Uncertain Alternatives	Draw from an uncertain urn	
	<i>Red (from 0.2 to 0.8)</i>	<i>Black (from 0.2 to 0.8)</i>
<i>Equality under Uncertainty</i>		
Ayan	80	30
Bashir	80	30
<i>Equal Uncertainty, Unequal Final Well-being</i>		
Ayan	80	30
Bashir	30	80

In what follows, for simplicity, I shall evaluate such alternatives while setting aside all considerations besides the well-being of the individuals in question. In a choice between *Inequality under Certainty*, which cures Ayan and leaves Bashir severely debilitated, and *Equal Risk, Unequal Final Well-being*, which will cure precisely one of them but gives each

an equal chance at being cured, the outlined egalitarian view chooses the latter. For there is, it holds, less unfairness when each is given an equal chance at a cure than when one is given a cure outright and the other has no chance at receiving it. *Equality under Certainty*, is, naturally, better still, since, by giving both a partially effective treatment that leaves them with a moderately good life with a level of well-being precisely midway between the grave disability and a full cure, it eliminates inequality in final well-being at no cost in total expected well-being. Finally, the view regards *Equality under Certainty* as just as good as *Equality under Risk*, because the latter also involves no inequality and offers each person the same expected well-being as the former.

This series of judgments is the upshot of combining a concern for eliminating unfair disadvantage in the value of prospects and in how people end up with a decision theory that is risk-neutral in personal good and moral value. It also has a grounding in a central idea of much of the post-WWII egalitarian literature on distributive justice, namely that an individual's life has a unity that a bare mass of people does not and that, consequently, moral precepts for pure intrapersonal trade-offs without inequality differ from precepts for trade-offs between distinct individuals' interests (Gauthier 1963, p. 121-7; Nagel 1970, p. 138; Rawls 1999, pp. 23-4). The integrality of a person's life gives us reason to make pure intrapersonal trade-offs with the aim of maximizing this person's expected well-being, as prudence dictates (given the assumed measure of well-being). The distinctions between persons, meanwhile, demands that when people's interest conflict, we have greater concern for those who are less well off. This idea motivates choosing *Equal Risk, Unequal Final Well-being* over *Inequality under Certainty*. For this choice involves an opposition of interests in having valuable prospects between Ayan and Bashir, which—given the fact that total expected well-being is constant—it resolves by maximizing the prospects of the least

well-off. It also motivates choosing *Equality under Certainty* over *Equal Risk, Unequal Final Well-being*, since this choice involves a conflict of final-well-being interests, with it being in the final-well-being interest of whoever would end up worse off under *Equal Risk, Unequal Final Well-being* that *Equality under Certainty* had been chosen instead, while the opposite would be in the final-well-being interest of whoever would end up better off under *Equal Risk, Unequal Final Well-being*. Given that these alternatives yield the same total well-being, the separateness of persons requires that this conflict be resolved in favor of the least well-off. Finally, respect for the unity of the individual supports regarding *Equality under Certainty* and *Equality under Risk* as equally choiceworthy, since the choice between them involves no conflicts of interest in terms of prospects or final well-being (and no inequality), so that we may choose any policy that maximizes each person's prospects.

2. Caution under Uncertainty

Let us now consider cases of uncertainty. Suppose again that Ayan and Bashir will suffer the aforementioned grave disability unless they are treated. You must either allocate resources towards the treatment described by *Equality under Risk* from Table 1, which, due to its extensive track record, you rationally believe offers a 0.5 chance of fully curing both and a 0.5 chance of being wholly ineffective for both, or instead provide a new treatment, *Equality under Uncertainty*, which will also either fully cure both or be wholly ineffective for both and for which the limited evidence available suggests that its chance of yielding a full cure ranges somewhere from 0.2 to 0.8. It is depicted in the lower part of Table 1. In line with Ellsberg's (1961) presentation, uncertainty is here represented by a random draw from an urn known to contain precisely 100 balls, all of which are either red or black, with the only information available being that at least 20 and at most 80 of these balls are red (i.e., no information is

available about the process by which the urn has been filled). Which treatment(s) is it permissible for you to provide?

In this choice, I submit that it is permissible for you to provide the treatment to which you can assign precise probabilities. Moreover, it would be permissible for you to have a strict preference for this treatment and to provide it even if it carried some small cost, in the sense of slightly worsening the final well-being outcomes for Ayan and Bashir. (To be precise, suppose that choosing *Equality under Risk* would result in cost c for each person in every event, so that if *Red* were drawn, Ayan and Bashir would each end up with a well-being of $80 - c$ and if *Black* were drawn, they would each end up with $30 - c$. My claim is that there is a $c > 0$ for which it would be permissible to choose *Equality under Risk*.)

The argument for this judgment proceeds in three steps (Joyce 2005, pp. 168-171; Gilboa, Postlethwaite, and Schmeidler 2009). First, rationality does not require you to go beyond the evidence and assign, arbitrarily, precise probabilities to the outcomes of each alternative that you might choose. Instead, it permits you simply to represent your beliefs in terms of ranges of probabilities assigned to each possible outcome as, say, the IPCC does for the “medium emissions” policy mentioned in the Introduction, when they judge that there is between 66% and 100% chance that this policy would lead to warming of more than 2.0 degrees. In our novel treatment example, this means you need not move beyond the assumption that the chance of this treatment fully curing Ayan and Bashir ranges from 20% to 80%.

Second, when you have only such imprecise probabilities, you cannot compute a single expected value for a prospect. But you can compute a range of such expected values. In the IPCC example, assuming that more warming is worse, the worst expected value of the

“medium emissions” policy will be one in which there is a 100% chance that it leads to more than 2.0 degrees warming and the best expected value of this policy is that there is only a 66% chance that it leads to such warming. All the IPCC’s information allows you to say is that the expected value of this medium emissions prospect is in the range given by these values. In our novel treatment example, this means that for each person, *Equality under Uncertainty* has an expected value in the range of 40 (the possible outcomes weighted by the least favorable probability distribution consistent with your evidence, that is, $0.8 \times 30 + 0.2 \times 80$) to 70 (the possible outcomes weighted by the most favorable probability distribution consistent with your evidence, that is, $0.2 \times 30 + 0.8 \times 80$).

Third, in the face of this range of expected values, it is permissible to be cautious, in the sense that in making an overall assessment of the uncertain prospect’s value, you may permissibly give more decision weight to the less good expected values than the better expected values. To apply it to our examples: when assessing the prospect associated with a policy of “medium emissions,” you are permitted to give more decision weight to the possibility that this would certainly lead to more than 2.0 degrees warming than to the possibility that this would only have a 66% chance of leading to such warming. And in the novel treatment case, you can permissibly take the prospective value of the novel medicine for each of Ayan and Bashir to be less than the mid-point between 40 and 70. (So less than 55, the expected value of the well-known medicine.)

The basic ideas in this argument are simple and attractive: no requirement to go beyond the evidence and a permission to be cautious in the face of lack of evidence. The upshot is that what is known as *uncertainty aversion* (a strict preference for prospects with precise probabilities over otherwise analogous prospects without such probabilities) is permissible.

There is a further dimension to the issue which arises for policymakers, namely, the attitudes towards uncertainty of the people whose prospects and fates hang in the balance. In general, respect for citizens' reasonable judgments of their own good makes it fitting for a policymaker to, as far as possible, track the rationally permissible attitudes of their citizens towards their own interests (Arneson 1990). Here, I include people's rationally permissible attitudes towards uncertainty in these judgments that policymakers have reason to respect. I also assume what I take to be a common situation for policymakers, which is that they do not know the uncertainty attitudes of every member of their population but do know the general social-scientific findings about these attitudes. Empirical studies suggest that in self-interested choices, both uncertainty aversion and uncertainty neutrality (which involves indifference between uncertain and analogous risky alternatives) are common, and uncertainty-loving behavior (which involves a strict preference for uncertain over comparable risky alternatives) is rare (Trautmann and van de Kuilen 2015, Table 1; Voorhoeve et al. 2016; Chew et al. 2018). It is therefore reasonable to hold that assuming a modest degree of uncertainty aversion on behalf of citizens would be the upshot of a procedure that minimized a reasonable measure of "aggregate distance" between citizens' diverse attitudes towards uncertainty. (For example, the mean attitude towards uncertainty suggested by the aforementioned studies would be one of modest uncertainty aversion.) If, for the reasons just outlined, uncertainty aversion is rationally permissible (possibly alongside uncertainty-neutral and uncertainty-loving behavior), this would therefore make it reasonable for the decision-maker to employ a degree of uncertainty aversion as a good approximation of (or reasonable compromise between) the differing reasonable attitudes of the individuals on whose behalf they are deciding.

It is important to note that despite its appeal, the rationality of uncertainty aversion is disputed. The reason is that on the assumption that uncertainty aversion is in tension with a core axiom of decision theory, the Sure Thing Principle. This means that uncertainty aversion has some unappealing implications.³

I cannot here review the extensive debate on the rational permissibility of uncertainty aversion. I will therefore briefly report my perspective on it, which is that the arguments show that not all independently attractive principles of rationality can be reconciled. In particular, there is at least an apparent tension between (i) the ideas that rationality does not require a decision-maker to posit precise probabilities for which they lack adequate ground and that a decision-maker is allowed a degree of caution in the face of such imprecision, and (ii) the idea that a decision-maker should respect other attractive principles of rational choice, such as the Sure Thing Principle. There are different reasonable ways of navigating this inconsistency, among which are uncertainty-averse decision principles (Gilboa et al. 2009; Siniscalchi 2009; Heal and Millner 2018).

There are several leading uncertainty-averse decision criteria. For concreteness, I shall here use a well-known, simple criterion which is often traced back to Leonard Hurwicz' work on decision making under ignorance (Hurwicz 1951). My conclusions hold for other leading criteria, e.g. those advanced in Gilboa and Schmeidler (1989) and Klibanoff et al. (2005).

On what is known as α -maxmin expected utility, the decision-maker values a prospect by taking $\alpha \times$ the worst expected value that is consistent with their information and priors and adding $(1 - \alpha) \times$ the best expected value that is so consistent, where $0 \leq \alpha \leq 1$ is the decision weight given to the worst expected value (Binmore 2009; Wakker 2010, sec. 11.5.) A

cautious evaluator will give more decision weight to the worst expected value—i.e., will have $\alpha > 0.5$ —and this will lead them to be uncertainty averse. An uncertainty-neutral evaluator will give equal weight to both—i.e., will have $\alpha = 0.5$. An uncertainty-loving evaluator will give greater weight to the best expected value, i.e., will have $\alpha < 0.5$. I shall in the rest of this chapter explore what follows if we assume a fixed, moderate degree of uncertainty aversion for all objects of evaluation—that is, both for the evaluation of a prospect for a single individual and a social prospect. This implies an invariant α somewhat larger than 0.5.

To illustrate, consider again the novel treatment represented by *Equality under Uncertainty*, and suppose for concreteness that $\alpha = 0.6$. The α -maxmin expected utility criterion then evaluates Ayan's uncertain prospect as follows: $0.6 \times (0.8 \times 30 + 0.2 \times 80)$ (the worst expected value) + $0.4 \times (0.2 \times 30 + 0.8 \times 80)$ (the best expected value) = 52. This is 3 units of well-being less than the corresponding risky treatment. In what follows, shall refer to this diminution of the value of an individual's prospects due to uncertainty as the "individual-level burden" of uncertainty. Naturally, the value of Bashir's prospects under *Equality under Uncertainty* is similarly depressed. But besides these individual-level burdens, this case involves what I shall call "social-level uncertainty" about the distribution of final well-being. For the facts that either both will end up fully cured or both will end up with a severe disability and that there are no precise probabilities for these outcomes may depress the prospective value of the social distribution of final well-being as compared to a counterpart policy under risk.

3. Cautious Egalitarianism

I shall now review a few key implications of incorporating this form of uncertainty aversion in the form of pluralist egalitarianism outlined in Section 1. In each instance, I shall connect the findings from simple cases to general considerations of justice and policymaking.

First, uncertainty-averse egalitarianism posits a novel object of egalitarian concern: the degree to which individual-level uncertainty depresses the value of individuals' prospects (Rowe and Voorhoeve 2018, pp. 255-6). An illustrative case of unequal burdens of uncertainty arises in the comparison of the uncertainty faced by people who live in regions and work in professions that are unlikely to be gravely disrupted by temperature rises (e.g., office workers in temperate zones) and those whose location and jobs are such that their lives and livelihoods would be severely affected by changes in the climate (e.g., farmers in marginal lands in the Sahel) (Denning et al. 2015). Another illustration concerns the differential burden of uncertainty around the negative impact on well-being of lockdowns to deal with COVID-19 in many countries. When first implemented, the range of potential impacts on well-being (and therefore the depressing effect of uncertainty on individuals' prospects at the moment of implementation) was arguably less for those in rich nations who could work remotely and who had access to government support if they should need it than it was for those in poorer nations who rely on the informal economy, whose work might be most disrupted by these measures, and who often have difficulty accessing social safety nets (Ray and Subramanian 2020). An uncertainty-averse, egalitarian view sees strong reasons to improve the prospects of those who face greater uncertainty, for example by insuring them

against the downside of their imprecisely estimated risks or by gaining information and thereby narrowing the imprecision in these estimates.

Second, the proposed cautious, egalitarian approach will favor policies for which a better basis is available for assigning probabilities to outcomes. This was already clear in the comparison made (at the end of Section 2) of *Equality under Risk* with its uncertain counterpart, *Equality under Uncertainty*. The same is true in a choice between *Equal Risk, Unequal Final Well-being* and its uncertain counterpart, *Equal Uncertainty, Unequal Final Well-being*, which is depicted in the lower part of Table 1. After all, the individual-level uncertainty created by the latter policy depresses the value of both individuals' prospects compared to its risky counterpart. One practical implication is that it is permissible for governments to go to greater expense to mitigate severely uncertain threats to life (e.g., posed by a novel pandemic) than to mitigate threats to which they can readily attach probabilities (e.g., posed by traffic accidents).⁴

A third key implication is that when, under uncertainty, some will gain and others will lose, uncertainty aversion reinforces the egalitarian tendency to allocate resources to those whoever will end up less well off (Rowe and Voorhoeve 2018, pp. 257-9). To see why, suppose that a policy-maker must choose between *Equal Uncertainty, Unequal Final Well-being* and *Equality under Certainty* in Table 1, but with the latter modified so that it comes at a cost c to each person's final (and prospective) well-being, so that it would yield only $55 - c$ for each person for sure, with $0 \leq c \leq 25$. On the proposed egalitarian view, in choosing between these options, we should consider both the value of individual prospects and the prospective value of the distribution of final well-being. Under *Equal Uncertainty, Unequal Final Well-being*, the value of individuals' prospects is depressed by the limited information

available about their likelihood of ending up badly off. An uncertainty-averse policymaker should therefore be willing to incur at least some small cost to eliminate this uncertainty.

Considering the distribution of final well-being, a downside of *Equal Uncertainty, Unequal Final Well-being* is inequality in how people will end up faring. Inequality aversion will therefore prompt a policymaker to incur a cost to eliminate this inequality.

It follows that both uncertainty aversion and inequality aversion will direct us to incur a cost to remove inequality in this kind of case. What is more, together, they will direct us to incur a higher cost to achieve equality than either one of these considerations alone would countenance (Rowe and Voorhoeve 2018, pp. 258-9). To understand why, assume for a moment that our inequality-averse policymaker was uncertainty neutral (that is, their $\alpha = 0.5$). They would then evaluate *Equal Uncertainty, Unequal Final Well-being* as equivalent to *Equal Risk, Unequal Final Well-being*. Suppose that in a choice between *Equal Risk, Unequal Final Well-being* and *Equality under Certainty* with a cost c to each person, the correct degree of inequality aversion will direct us to incur a cost of up to, but no greater than, c^* units of prospective well-being for each person, so that both *Equal Uncertainty, Unequal Final Well-being* and *Equal Risk, Unequal Final Well-being* would be equivalent to giving each of Ayan and Bashir $55 - c^*$ for sure. Next, assume that our policymaker becomes uncertainty averse (that is, their $\alpha > 0.5$). They will then find *Equal Uncertainty, Unequal Final Well-being* strictly worse than *Equal Risk, Unequal Final Well-being*. By transitivity, they will then regard the uncertain alternative as strictly worse than giving each of Ayan and Bashir $55 - c^*$. In other words, they will find *Equal Uncertainty, Unequal Final Well-being* to be as good as *Equality under Certainty* only for a cost larger than c^* . We can conclude that in cases in which individual-level uncertainty will lead to some faring better than others, uncertainty-averse egalitarianism justifies incurring a larger cost in order to achieve both

equality and certainty than an uncertainty-neutral egalitarian view would countenance. After all, under these circumstances, directing benefits from the lucky to the unlucky reduces the stakes for each and thereby reduces the burden of uncertainty; naturally, it also diminishes inequality. A policy issue to which this may be relevant is levying “windfall taxes” on firms and people who gain due to severely uncertain economic developments, with these taxes spent on improving the situation of the losers. (This is a policy that has been considered in the UK, for example, in response to the COVID-19 crisis, see Cowburn 2021.) Such policies that reduce the variability of incomes under uncertainty will be valuable both because they reduce the burden of uncertainty and because they reduce inequality.

A fourth key conclusion starts from the observation that individual-level uncertainty need not imply social-level uncertainty about the distribution of final well-being. To see why, consider again *Equal Uncertainty, Unequal Final Well-being* in Table 1. The individual-level uncertainty depresses the value of everyone’s prospects. However, it has no social-level uncertainty about the value of the possible distributions of final well-being. For the anonymized distribution of final well-being is known: one person will be fully cured, another will remain severely disabled.

This has an important implication for how the ranking of policies under risk compares to the ranking of their counterpart policies under uncertainty. *Equality under Risk* is, on the proposed egalitarian view, strictly preferred to *Equal Risk, Unequal Final Well-being*, because the former eliminates all inequality without any loss in terms of the value of individual prospects or in the value of the prospective distribution of final well-being. However, the ranking of these policies’ uncertain counterparts is less straightforward. In terms of the value of individual prospects, *Equality under Uncertainty* and *Equal*

Uncertainty, Unequal Final Well-being are identical. However, in terms of the prospective value of the possible distributions of final well-being, a concern for equality and a concern to reduce uncertainty may pull in opposite directions. On the one hand, it counts in favor of *Equality under Uncertainty* that it eliminates all inequality. On the other hand, because under this policy, everyone sinks or swims together, it generates problematic uncertainty at the collective level. In contrast, *Equal Uncertainty, Unequal Final Well-being* does not generate such collective-level uncertainty. This is one respect in which *Equal Uncertainty, Unequal Final Well-being* may be better.⁵

The view I put forward here does not pronounce on which of these two policies is superior overall. The key conclusion is just that uncertainty aversion may oppose the solidaristic, egalitarian impulse to bind everyone's fates together. In doing so, it changes egalitarianism in one important way (Rowe and Voorhoeve 2018, pp. 261-2). Under risk, the outlined egalitarian view has a tendency to allocate benefits away from the lucky and towards the unlucky *if and only if the lucky are (or would be without an egalitarian allocation) better off than others and the unlucky are (or would be) less well off than others*. By way of illustration, as we have seen, the proposed form of egalitarianism is indifferent between *Equality under Risk* and *Equality under Certainty* in Table 1. Moreover, if *Equality under Certainty* could be purchased only at a small cost c to each person's well-being (so that it would yield only $55 - c$ for each person for sure, with c positive and small), the proposed egalitarian view would be unwilling to pay any cost in order to redistribute from the better off potential futures of Ayan and Bashir to their less well off potential futures. By contrast, under uncertainty, a cautious egalitarian will see reason to direct benefits from the lucky towards the unlucky *even when these are merely two potential futures of the same person and there is no inequality*. To see this, compare *Equality under Uncertainty* with *Equality under*

Certainty. Due to uncertainty aversion, the latter is clearly preferable. Moreover, if *Equality under Certainty* could be purchased only at a cost c to each person's final (and prospective) well-being (so that it would yield only $55 - c$ for each person for sure), the proposed view would strictly prefer *Equality under Certainty* over *Equality under Uncertainty* for some small, positive c . Under uncertainty, cautious egalitarianism is therefore keen to direct benefits away from Ayan and Bashir's better possible futures towards their worse possible futures, even when their rosier futures would not involve them being better off than others.

As a practical matter, it follows from the proposed view that governments have special reason to make provision for collective setbacks to which they are not able to assign a precise probability. As a concrete illustration, if natural resource revenues are of this kind, then for governments that are highly dependent on revenues from these resources, this favors instruments such as hedges against price falls (used, for example, by Mexico to cover its oil revenues, see Reuters 2015) or creating fiscal space to support the economy in the face of a price collapse (as has been practiced in Chile, see Céspedes et al. 2014).

4. Conclusion

I have argued that in uncertain situations, it is permissible for a policymaker to consider a range of expected values for each policy, rather than a single expected value. I have also argued that, in response to this range, it is permissible for a policymaker to assign greater decision weight to the lower expected values within this range. One reason I have offered for this approach is that a substantial share of the population whose fates are at stake in their decisions are likely to be uncertainty averse and very few are likely to be uncertainty loving,

so that a modest degree of uncertainty aversion in public decision making is a reasonable compromise.

I also explored some key implications of incorporating such uncertainty aversion into a pluralistic egalitarian theory of justice and for using such a theory for policy evaluation. I argued that uncertainty aversion reinforces egalitarian reasons to reduce unfair inequalities and to resolve *interpersonal* conflicts of interest in favor of the less well off. Moreover, it gives us new reasons to make *intrapersonal* trade-offs under uncertainty in a way that favors the person's less fortunate potential futures. The upshot is a theory of justice which offers stronger reasons for safety net policies such as universal health coverage, unemployment insurance, and disability pay that guard against individual and collective misfortune. Such policies are commonly defended as valuable because they improve people's prospects by reducing risks in relation to income and health in an efficient manner and because they reduce inequalities (Barr 2012, WHO 2014). But a consideration of severely uncertain situations reveals further functions of such a safety net. By aiding the unfortunate, it both reduces the depressing impact of uncertainty on the value of individuals' prospects and reduces policymakers' uncertainty about how social outcomes. This finding is relevant for our two opening examples. In the context of climate change, higher emissions pathways are associated with greater variability in the moral value of possible outcomes and therefore with a larger disvalue of uncertainty (Millner et al. 2013); they are also projected to generate greater inequality (Denning et al. 2015). A cautious, inequality-averse approach will therefore hold that we have strong reasons to lower emissions (and more reasons than a common, expected value maximizing utilitarian approach would register). In the context of the COVID-19 pandemic, this approach reinforces reasons to develop treatments and vaccines (since these will tend to improve worse possible futures) and to introduce social and

movement measures to contain the spread of disease so long as these are accompanied by income support to the worst off (Adler et al. 2020; Ray and Subramanian 2020). In short, uncertainty adds to policymakers' reasons to make provision for the least fortunate.

References

- Adler, M. (2019) *Measuring Social Welfare: An Introduction*. Oxford: Oxford University Press.
- Adler, M., R. Bradley, M. Ferranna, M. Fleurbaey, J. Hammitt, and A. Voorhoeve. (2020) *How to Assess the Well-being Impacts the COVID-19 Pandemic and Three Policy Types: Suppression, Control and Uncontrolled Spread*. Policy brief published as part of the final communique of the T20 (Thinktank 20) accompanying the G20 in Saudi Arabia. https://t20saudiarabia.org.sa/en/briefs/Pages/Policy-Brief.aspx?pb=TF4_PB8
- Al-Najjar, N. and J. Weinstein. (2009) "The Ambiguity Aversion Literature: A Critical Assessment," *Economics and Philosophy* 25: 249-84.
- Arneson, R. (1990) "Liberalism, Distributive Subjectivism, and Equal Opportunity for Welfare." *Philosophy & Public Affairs*, 19(2): 158-94.
- Arneson, R. (1997) "Postscript to 'Equality and Equal Opportunity for Welfare'," in L. Pojman and R. Westmoreland (ed.) *Equality: Selected Readings*, Oxford: Oxford University Press, pp. 238-41.
- Avery, C., W. Bossert, A. Clark, G. Ellison, S. Fisher Ellison. (2020) "Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists." *NBER Working Paper 27007, version of April 2020*.
<http://www.nber.org/papers/w27007>
- Barr, N. (2012) *The Economics of the Welfare State, 5th edition*. Oxford: Oxford University Press.

- Berger, L., N. Berger, V. Bosetti, I. Gilboa, L. P. Hansen, C. Jarvis, M. Marinacci, R. D. Smith. (2020) “Uncertainty and Decision making During a Crisis: How to Make Policy Decisions in the COVID-19 Context?” *Innocenzo Gasparini Institute for Economic Research Working Paper n. 666*, July 2020 version.
- Binmore, K. (2009) *Rational Decisions*. Princeton: Princeton University Press.
- Bognar, G. and I. Hirose. (2014) *The Ethics of Health Care Rationing: An Introduction*. Abingdon: Routledge.
- Bradley, R. (2017) *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.
- Broome, J. (1990) “Fairness,” *Proceedings of the Aristotelian Society* 91: 87–101.
- Céspedes, L. F., E. Parrado, and A. Velasco. (2014) “Fiscal Rules and the Management of Natural Resource Revenues: The Case of Chile,” *Annual Review of Resource Economics* 6: 105-132.
- Chew, S. H., M. Ratchford, J. S. Sagi. (2018) “You Need to Recognise Ambiguity to Avoid It,” *The Economic Journal* 128 (614): 2480-2506.
- Cowburn, A. (2021) “Covid: Government ‘Considering Excess Profits Tax’ on Companies Cashing in During Pandemic,” *The Independent*
<https://www.independent.co.uk/news/uk/politics/covid-profits-amazon-rishi-sunak-budget-b1798839.html> [accessed 28 February 2021].
- Denning, F., M. B. Budolfson, M. Fleurbaey, A. Siebert, and R. H. Socolow. (2015) “Inequality, Climate Impacts on the Future Poor, and Carbon Prices,” *Proceedings of the National Academy of Sciences* 120 (52): 15827–32.
www.pnas.org/cgi/doi/10.1073/pnas.1513967112

- Dietz, S. (2014) “Climate Change Mitigation as Catastrophic Risk Management,”
Environment: Science and Policy for Sustainable Development, 56:6, 28-36, DOI:
 10.1080/00139157.2014.964096
- Ellsberg, D. (1961) “Risk, Ambiguity, and the Savage Axioms,” *The Quarterly Journal of Economics* 75 (1961): 643—69.
- Fleurbaey, M. (2010) “Assessing Risky Social Situations,” *Journal of Political Economy*,
 118: 649-80.
- Frigg, R. S. Bradley, H. Du, and L. A. Smith. (2014) “Laplace’s Demon and the Adventures of
 His Apprentices,” *Philosophy of Science*, 81 (1): 31-59.
- Gauthier, D. (1963) *Practical Reasoning*. Oxford: Clarendon Press.
- Gilboa, I. and D. Schmeidler. (1989) “Maximin Expected Utility with Non-Unique Prior,”
Journal of Mathematical Economics 18 (1989): 141—53.
- Gilboa, I., A. Postlethwaite and D. Schmeidler. (2009) “Is It Always Rational to Satisfy
 Savage’s Axioms?” *Economics and Philosophy* 25: 285–96.
- Hammit, J. (2020) “Valuing mortality risk in the time of COVID-19.” *Journal of Risk and
 Uncertainty* 61: 129–54.
- Heal, G. and A. Millner. (2018) “Uncertainty and Decision making in Environmental
 Economics: Conceptual Issues”, in P. Dasgputa, S. Pattanayak, and V.K. Smith (eds.),
Handbook of Environmental Economics (Vol 4), Elsevier: North-Holland, chap. 10.
- Hurwicz, L. (1951). “Optimality Criteria for Decision Making under Ignorance,” *Cowles
 Commission Discussion Paper, Statistics*, 370.
- Inoue, A. and K. Miyagishima. (2021) “A Defense of Pluralist Egalitarianism under Severe
 Uncertainty: Axiomatic Characterization,” *manuscript*, version of 9/03/2021.

- IPCC. (2014) *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC, Geneva.
- Joyce, J. M. (2005) “How Probabilities Reflect Evidence,” *Philosophical Perspectives: Epistemology* 19: 153—78.
- Keynes, J. M. (1937) “The General Theory of Employment,” *The Quarterly Journal of Economics*, 51 (2): 209–23.
- Klibanoff, P., M. Marinacci, and S. Mukerji. (2005) “A Smooth Model of Decision Making under Ambiguity,” *Econometrica* 73: 1849–92.
- Knight, F. (1921) *Risk, Uncertainty and Profit*. Boston, MA: Hart, Schaffner and Marx.
- Liu, J.-T., J. K. Hammitt, J.-D. Wang, and M.-W. Tsou. (2005) “Valuation of the Risk of SARS in Taiwan,” *Health Economics* 14: 83-91.
- Miles, D., M. Stedman, and A. Heald. (2020) “Living with COVID-19: Balancing Costs against Benefits in the Face of the Virus,” *National Institute Economic Review* 253: R60-R76. doi:10.1017/nie.2020.30.
- Millner, A., S. Dietz and A. Heal. (2013) “Scientific Ambiguity and Climate Policy,” *Environmental and Resource Economics* 55:21–46
DOI 10.1007/s10640-012-9612-0
- Nagel, T. (1970) *The Possibility of Altruism*. Princeton: Princeton University Press.
- Otsuka, M. and A. Voorhoeve. (2018) “Equality versus Priority,” in S. Olsaretti (ed.), *The Oxford Handbook of Distributive Justice*, Oxford: Oxford University Press, pp. 65-85.
- Ray, D. S. and S. Subramanian. (2020) “India's Lockdown: An Interim Report,” *Indian Economic Review*, online first doi: 10.1007/s41775-020-00094-2.
- Rawls, J. (1999) *A Theory of Justice, rev. ed.* Cambridge, M.A.: Harvard University Press.

- Reuters. (2015) “Mexico Wraps \$1.1 Billion Oil Options Hedge to Lock in \$49 Floor,” August 20, 2015. <http://www.reuters.com/article/us-mexico-oil-idUSKCN0QP0X020150820> [Accessed 22 February 2021].
- Rowe, T. and A. Voorhoeve. (2018) “Egalitarianism under Severe Uncertainty,” *Philosophy & Public Affairs* 46 (3): 239-68.
- Siniscalchi, M. (2009) “Two Out of Three Ain’t Bad: A Comment on ‘The Ambiguity Aversion Literature: A Critical Assessment’,” *Economics and Philosophy* 25: 335—56.
- Stefansson, H. O. (this volume) “The Economics and Philosophy of Risk,” in C. Heilmann and J. Reiss (ed.) *The Routledge Handbook of Philosophy of Economics*. Abingdon-on-Thames: Routledge.
- Temkin, L. (2001) “Inequality: A Complex, Individualistic, and Comparative Notion,” *Philosophical Issues*, 11: 327-53.
- Treich, N., (2010) “The Value of a Statistical Life under Ambiguity Aversion,” *Journal of Environmental Economics and Management* 59: 15–26.
- Trautmann, S. and G. van de Kuilen (2015) “Ambiguity Attitudes,” in G. Keren and G. Wu (ed.) *The Wiley Blackwell Handbook of Judgment and Decision Making*, Chichester: Wiley, pp. 89–116.
- Ulph, A. (1982) “The Role of Ex Ante and Ex Post Decisions in the Valuation of a Life,” *Journal of Public Economics*, 18: 265-76.
- Voorhoeve, A. and M. Fleurbaey. (2012) “Egalitarianism and the Separateness of Persons,” *Utilitas*, 24: 381-98.
- Voorhoeve, A., K. Binmore, A. Stefansson and L. Stewart. (2016) “Ambiguity Attitudes, Framing and Consistency,” *Theory and Decision*, 81: 313-37.

Voorhoeve, A. (2021) “Equality for Prospective People: A Novel Statement and Defence.”

Utilitas, in press.

Wakker, P. (2010) *Prospect Theory for Risk and Uncertainty*. Cambridge: Cambridge

University Press.

World Health Organization. (2014) *Making Fair Choices on the Path to Universal Health*

Coverage: Final report of the WHO Consultative Group on Equity and Universal

Health Coverage. WHO: Geneva.

Acknowledgements

I am grateful to Richard Bradley, Roman Frigg, Kaname Miyagishima, Joe Roussos, and Thomas Rowe for comments. Work on this article was supported through the Bergen Centre for Ethics and Priority Setting’s project ‘Decision Support for Universal Health Coverage’, funded by NORAD grant RAF-18/0009.

Biographical note

Alex Voorhoeve is Professor in the Department of Philosophy, Logic and Scientific Method at the London School of Economics and Political Science and Visiting Professor of Ethics and Economics at Erasmus University Rotterdam. He has published widely on the theory and practice of distributive justice, especially as applied to health. He is author of *Conversations on Ethics* (Oxford University Press, 2009), a collection of dialogues with leading philosophers and economists and co-author of *Making Fair Choices on the Path to Universal Health Coverage* (World Health Organization, 2014).

¹ In moral and political philosophy, an exception to the neglect of severe uncertainty has been the discussion of how to make decisions behind John Rawls's veil of ignorance, which creates a severely uncertain situation by denying people knowledge of the probability of ending up in any particular social position (Rawls 1999, p. 134). In welfare economics, contributions that take account of this aspect of severe uncertainty in policy evaluation focus on environmental policy and pandemics. See, e.g., Liu et al. (2005), Treich (2010), Heal and Millner (2018), Berger et al. (2020), and Inoue and Miyagishima (2021).

² This formulation leaves open the precise way these different aims of reducing inequality and promoting well-being are defined. For concreteness, I shall assume that the degrees to which prospects and final well-being are promoted are given by their total value. See Fleurbaey (2010); Voorhoeve and Fleurbaey (2016) and Voorhoeve (2021, appendix) for a proposed more precise formulation of pluralist egalitarianism. The proposed form of egalitarianism builds on an extensive literature which emphasizes the importance of both people's prospects and final well-being, including Ulph (1982), Cohen (1989), Broome (1990), and Temkin (1993).

³ For a discussion on the Sure Thing Principle and uncertainty aversion, see the chapter by Stefansson in this volume. For discussion of the problems to which this violation gives rise, see Al-Najjar and Weinstein (2009).

⁴ Such greater expense to prevent imprecise risks aligns with surveys on the value of reductions in fatality risks, in which individuals tend to place a premium on reductions in imprecisely over precisely specified chances of death (see Hammitt 2020, pp. 140-8).

⁵ This is true, at least, for some policymakers with a modest degree of inequality aversion. An extremely inequality averse policymaker may well hold that there is no respect in which *Equal Uncertainty, Unequal Final Well-being* is more valuable. For example, take a policymaker who uses the maximin rule to evaluate distributions of final well-being. They

will hold that *Equality under Uncertainty* has a better distribution of final well-being than *Equal Uncertainty, Unequal Final Well-being* in one state of the world and an equivalent distribution of final well-being in another, so that the former dominates the latter, despite uncertainty.