

I. MORAL DECISIONS UNDER SEVERE UNCERTAINTY

In the spring of 2009, a novel strain of the H1N1 influenza virus, containing a never before witnessed combination of gene segments from human influenza, two forms of swine influenza, and avian influenza,¹ was declared a global pandemic. The UK Government had to decide whether to undertake, at a cost of £1.2 billion (USD 1.9 billion at the time, equivalent to 1 percent of that year's health budget), an extensive set of preparatory measures, including the purchase of both antiviral medication and a novel vaccine in quantities sufficient to cover the entire UK population, or whether instead to take substantially less costly measures, which would involve having only a limited supply of these medicines and vaccines at hand.² The possible

This article was presented at Erasmus University Rotterdam; George Washington University; the LSE; the universities of Cambridge, Edinburgh, Manchester, Minho, Montréal, Oxford, Pavia, Reading, Salzburg, and Warwick; and the U.S. National Institutes of Health. We thank our audiences and Matthew Adler, Gwyn Bevan, Ken Binmore, Richard Bradley, Campbell Brown, Susanne Burri, Nicolas Cote, Hilary Greaves, Keith Hyams, Anna Mahtani, Joseph Mazor, William MacAskill, Joseph Millum, Adam Oliver, Michael Otsuka, Tom Sorrell, Kai Spiekermann, Robert Steel, H. Orri Stefansson, Bastian Steuwer, Martin van Hees, Peter Wakker, James Wilson, and two anonymous referees for this journal for comments. Work on this article was supported by the British Arts and Humanities Research Council through grant AH/J006033/1. Alex Voorhoeve also thanks the National Institutes of Health for a Visiting Scholarship during which part of this article was completed. The authors alone are responsible for the views expressed in this article. They do not represent any position, decision, or policy of the U.S. National Institutes of Health, the U.S. Public Health Service, or the U.S. Department of Health and Human Services.

1. Food and Agriculture Organization, "The Human Influenza Due to a Novel Subtype H1N1," *Emergency Prevention Systems Watch* 1 (2009). <http://www.fao.org/3/a-ak061e.pdf>. Accessed August 22, 2018.

2. The information about the UK's decision-making during the H1N1 pandemic in this paragraph is based on Deirdre Hine, *An Independent Review of the UK Response to the 2009 Influenza Pandemic* (London: The Cabinet Office, 2010), and Adam Oliver, "Ambiguity Aversion and the UK Government's Response to Swine Flu," in *Behavioural Public Policy*, ed. Adam Oliver (Cambridge: Cambridge University Press, 2013), 16–31.

scenarios presented by civil servants to the government involved greatly varying degrees of disease spread and severity: the “reasonable range” of possibilities extended from a few hundred fatalities to 65,000 deaths, and government planning for a flu pandemic called for consideration of the possibility that there would be up to 750,000 deaths. (A normal seasonal flu kills between 2,000 and 4,000 people in the United Kingdom.) On professional estimates, the widespread use of antivirals and universal availability of the newly developed vaccines had the potential to significantly lower the death rate in the worse scenarios, but it was unclear by how much. Because of the limited data available about the severity and speed of spread of the virus as well as about the effectiveness of the novel vaccine and antivirals, expert advisors and key decision-makers believed they did not have sufficient grounds for assigning precise probabilities to the possible outcomes of the two principal policy alternatives in front of them. The UK Government was, therefore, in a situation of quite severe uncertainty, in the sense famously articulated by John Maynard Keynes:

By “uncertain” knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; (...). Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper (...). About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.³

Following seminal work by Daniel Ellsberg, this form of uncertainty is also commonly referred to in the economic literature as “ambiguity.”⁴ An uncertain (or ambiguous) situation so defined contrasts with a merely risky situation in which the decision-maker is in a position to assign precise probabilities to all relevant potential outcomes of the alternatives they have to choose between. We emphasize that in the senses in which

3. John Maynard Keynes, “The General Theory of Employment,” *The Quarterly Journal of Economics* 51 (1937): 209–23, at pp. 213–4. See also Frank Knight, *Risk, Uncertainty and Profit* (Boston, MA: Hart, Schaffner and Marx, 1921).

4. Daniel Ellsberg, “Risk, Ambiguity, and the Savage Axioms,” *The Quarterly Journal of Economics* 75 (1961): 643–69. See also Richard Bradley and Mareile Drechsler, “Types of Uncertainty,” *Erkenntnis* 79 (2014): 1225–48.

we use these terms, both “risk” and “uncertainty” are subjective notions—they pertain to the beliefs about the chances of all relevant outcomes of their decisions that a rational decision-maker is in a position to form on the basis of their prior beliefs and the evidence available to them.

Uncertain situations in this sense are common. For example, with regard to many policy-relevant possibilities, the International Panel on Climate Change (IPCC) reports only probability intervals, because the best available information does not suffice for the assignment of precise probabilities. That is, they make statements of the following kind:

It is *likely* [official translation: there is a chance between 0.66 and 1] that land temperatures over Africa will rise faster than the global land average, particularly in the more arid regions.⁵

Uncertainty also arises in private decisions, such as when a doctor is considering a novel drug for multiple sclerosis for their patient, and the evidence, combined with the decision-makers’ prior beliefs, does not permit them to nonarbitrarily assign precise probabilities to this treatment’s effects.⁶

Given the ubiquity of uncertainty, it is striking that the philosophical discussion of the requirements of egalitarian distributive justice under uncertainty is far less developed than it is for conditions of risk.⁷ Here, we

5. IPCC, “Africa,” in *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. Christopher Field, Vincente Barros, David Dokken, Katharine Mach, Michael Mastrandrea, T. Eren Bilir, Monalisa Chatterjee, Kristie Ebi, Yuka Otsuki Estrada, Robert Genova, Betelhem Girma, Eric Kissel, Andrew Levy, Sandy MacCracken, Patricia Mastrandrea, and Leslie White (Cambridge: Cambridge University Press), pp. 1199–243 at p. 1202. For discussion of the sources of severe uncertainty in climate science, see Geoffrey Heal and Antony Millner, “Uncertainty and Decision Making in Climate Change Economics,” *Review of Environmental Economics and Policy* 8 (2014): 120–37.

6. See, for example, Gustavo Saposnik, Angel Sempere, Daniel Prefasi, Daniel Selchen, Christian Ruff, Jorge Maurino, and Philippe Tobler, “Decision-making in Multiple Sclerosis: The Role of Aversion to Ambiguity for Therapeutic Inertia among Neurologists,” *Frontiers in Neurology* 8 (2017), article 65.

7. The literature on egalitarian distributive justice under risk is voluminous. For some recent contributions, see Marc Fleurbaey, “Assessing Risky Social Situations,” *Journal of Political Economy* 118 (2010): 649–80; Johann Frick, “Contractualism and Social Risk,” *Philosophy & Public Affairs* 43 (2015): 175–223; David McCarthy, “Distributive Equality,” *Mind* 124 (2015): 1045–109; and Alex Voorhoeve and Marc Fleurbaey, “Priority or Equality for Possible People?” *Ethics* 126 (2016): 929–54. By contrast, the philosophical literature on

take a step toward remedying this lack. We take as our point of departure a recent, pluralist egalitarian theory of distributive justice under risk. We propose and defend a novel extension of this view for uncertain situations and trace some key implications for policy decisions.

We proceed as follows. In Section II, we summarize the egalitarian view for decision-making under risk that we take as our point of departure. In Section III, we introduce a cautious or “uncertainty-averse” decision criterion that we will appeal to throughout. On this criterion, uncertainty represents a burden in the sense that it reduces the value of a prospect. In subsequent sections, we explore novel implications generated by the interplay of the twin aims of reducing the burden of uncertainty and limiting inequality. In Section IV, we discuss cases where these aims are congruent. We show that our view provides novel reasons to direct resources toward those who have worse prospects or outcomes than others. In Section V, we consider cases where uncertainty aversion and inequality aversion are in tension. We show that our view weakens the egalitarian impulse to ensure that everyone sinks or swims together, because it gives great weight to eliminating the possibility of collective misfortune. In Section VI, we provide a new perspective on the debate between utilitarians and egalitarians. We demonstrate that if aversion to uncertainty is permissible, then utilitarians cannot wield a favorite argument against

egalitarianism under uncertainty has focused rather narrowly on the question whether egalitarian principles can be derived from John Rawls’s veil of ignorance, which creates an uncertain situation by denying people knowledge of the probability of ending up in any particular social position. See Rawls, *A Theory of Justice*, rev. ed. (Cambridge, Mass.: Harvard University Press, 1999), p. 134; Leon Ellsworth, “Decision-Theoretic Analysis of Rawls’ Original Position,” in *Foundations and Applications of Decision Theory: Volume II: Epistemic and Social Applications*, ed. Clifford Alan Hooker, James Leach, and Edward McClennen (Dordrecht: Springer, 1978), pp. 29–45; Alexander Kaufman, “A Satisfactory Minimum Conception of Justice: Reconsidering Rawls’s Maximin Argument,” *Economics and Philosophy* 29 (2013): 349–69; and H. Orri Stefansson, “Ambiguity Aversion behind the Veil of Ignorance,” unpublished manuscript. For a rare exception to the veil of ignorance approach among philosophers, see Caspar Hare, “Risk and Radical Uncertainty in HIV Research,” *Journal of Medical Ethics* 43 (2017): 87–9.

Environmental welfare economists, however, have begun to discuss how to evaluate policies in the face of the severe uncertainties involved in climate change. See Antony Millner, Simon Dietz, and Geoffrey Heal, “Scientific Ambiguity and Climate Policy,” *Environmental and Resource Economics* 55 (2013): 21–46; and Heal and Millner, “Uncertainty and Ambiguity in Environmental Economics: Conceptual Issues,” unpublished manuscript on <http://personal.lse.ac.uk/MILLNER/Site/Home.html>. Accessed August 24, 2018.

egalitarians. We summarize our principal findings and their relevance for a range of policy decisions in Section VII, where we also return to our opening H1N1 influenza case.

Before proceeding, we emphasize that our aim is merely to propose an egalitarian view that incorporates a set of rationally and morally permissible (rather than required) differential attitudes toward risk and uncertainty. In the service of this aim, we assume orthodox decision theory under risk, because leading alternatives to the orthodoxy under risk see conformity with the orthodoxy as rationally permissible.⁸ However, we pair it with an unorthodox (although well known) decision principle under uncertainty, which yields the orthodoxy in the special case in which the decision-maker assigns precise probabilities to each outcome. The challenges to the orthodoxy posed by uncertainty are unique, because they involve a decision-maker lacking the grounds for forming reasoned beliefs of the kind that play a central role in making the orthodoxy plausible. It is therefore coherent, and indeed common, to endorse the orthodoxy in risky cases but not in cases of uncertainty.⁹ This approach also allows us to focus squarely on unexplored issues involved in confronting uncertainty.

II. EGALITARIANISM UNDER RISK

In this section, we describe the egalitarian view for decisions under risk that we aim to extend to uncertain cases. To proceed at pace to novel ideas in subsequent sections, we do not offer a full defense of this view, which is provided elsewhere.¹⁰ We shall refer to it as follows:

Pluralistic egalitarianism: We should aim to improve people's prospects for well-being, raise total well-being, and reduce inequality in both

8. See, for example, Lara Buchak, *Risk and Rationality* (Oxford: Oxford University Press, 2013).

9. See, for example, Gilboa and David Schmeidler, "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics* 18 (1989): 141-53; Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji, "A Smooth Model of Decision Making under Ambiguity," *Econometrica* 73 (2005): 1849-92; and Ken Binmore, *Rational Decisions* (Princeton: Princeton University Press, 2009).

10. See Fleurbaey, "Assessing Risky Social Situations;" Voorhoeve and Fleurbaey, "Priority or Equality;" and Michael Otsuka and Voorhoeve, "Equality versus Priority," in *The Oxford Handbook of Distributive Justice*, ed. S. Olsaretti (Oxford: Oxford University Press, 2018), pp. 65-85.

people's prospects and in their final well-being (how well their lives end up going).

With respect to each individual's fate, we will assume that we are concerned with the distribution of a cardinal, interpersonally comparable measure of lifetime well-being derived from idealized preferences satisfying the von Neumann–Morgenstern axioms under risk. On this measure, a prospect has higher expected well-being for a person just in case it would be preferred after rational and calm deliberation with all pertinent information while attending to that person's self-interest only.¹¹ One prospect has the same expected well-being as another prospect for a person just in case such deliberation would yield indifference between the two prospects.

To illustrate our egalitarian view under risk, imagine the following situation of a resource allocation manager in the National Health Service. Two 10-year-old children, Ann and Bea, have just been diagnosed with an illness which, if untreated, will leave them completely blind and with a lifetime well-being of 50 (a moderately good quality of life); if fully cured, each would have a lifetime well-being of 80 (a very good quality of life).¹² Both are strangers to the decision-maker and to each other. Unfortunately, the resources at the decision-maker's disposal do not suffice to fully cure both Ann and Bea for sure. Below, we will describe the alternatives open to them. To link up with Ellsberg's paradigmatic 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.¹³

11. This *measure* does not presuppose a specific view on what well-being *is*. One can hold that two alternatives yield the same expected well-being for a person precisely when they would, if ideally rational and wholly self-interested, be indifferent between these alternatives while also holding that well-being does not consist of preference satisfaction. For one might maintain that well-being consists of something other than preference satisfaction and that the specified idealized preferences fully track the size of this other thing. See Otsuka and Voorhoeve, "Why It Matters that Some Are Worse Off than Others: An Argument against the Priority View," *Philosophy & Public Affairs* 37 (2009): 171–99, at pp. 172–3, n. 3.

12. The numbers given for life-time well-being correspond roughly to those yielded by the well-known Health Utilities Index, Mark III (which is a von Neumann–Morgenstern measure of health-related well-being), for, respectively: living 10 years in full health followed by 70 years being completely unable to see; and living 80 years in full health. See *The Health Utilities Index, Mark III*. <http://www.healthutilities.com/hui3.htm>. Accessed July 25, 2017. On this scale, zero well-being is a life of equivalent value to the person to never having existed.

13. Ellsberg, "Risk, Ambiguity, and the Savage Axioms."

(Although it may seem odd to speak of providing treatments that are effective conditional on the draw of a ball of a particular color from an urn, this is merely a device to depict treatments for which the decision-maker rationally assigns precise probabilities to each possible outcome.) Table 1 lists the final well-being for Ann and Bea given each possible draw from this urn for each of the alternatives we will presently consider.¹⁴ We use red_r and $black_r$ to represent the possible draws from this risky urn, and p_{red} and p_{black} for the probability of these draws. Throughout, for simplicity, we will consider only Ann's and Bea's well-being; we will not consider how their level of well-being relates to that of further people.

Suppose first that the decision-maker must choose between only the following two alternatives:

Inequality under Certainty: Cure Ann and leave Bea to go wholly blind.

Equal Risk, Unequal Final Well-being: This treatment will either cure Ann and be entirely ineffective for Bea (leaving her to go wholly blind), or, instead, be entirely ineffective for Ann (leaving her to go wholly blind) and cure Bea. These results are equally likely.

Our pluralist egalitarian view requires choosing the latter. For it is concerned with limiting unfairness, and although both alternatives yield unfair inequality in final well-being, there is, on this view, less unfairness overall when each is given an equal shot at a cure than when one child is given a cure outright and the other has no chance at receiving it.¹⁵

Next, suppose that the decision-maker has a further alternative available:

Equality under Risk: this treatment will either cure both children, or be wholly ineffective for both, with each result being equally likely.

14. For discussion of a similar set of alternatives, see Luc Bovens, "Concerns for the Poorly Off in Ordering Risky Prospects," *Economics and Philosophy* 31 (2015): 397–429.

15. See, for example, Broome, "Uncertainty and Fairness;" Richard Arneson, "Postscript to 'Equality and Equal Opportunity for Welfare,'" in *Equality: Selected Readings*, ed. Louis Pojman and Robert Westmoreland (Oxford: Oxford University Press, 1997), pp. 238–41 at p. 240; and Larry Temkin, "Inequality: A Complex, Individualistic, and Comparative Notion," *Philosophical Issues* 11 (2001): 327–53, at pp. 337–9.

Table 1.

Alternative	Event	
	Draw from a risky urn	
	red_r ($p_{red} = 0.5$)	$black_r$ ($p_{black} = 0.5$)
<i>Inequality under Certainty</i>		
Ann	80	80
Bea	50	50
<i>Equal Risk, Unequal Final Well-being</i>		
Ann	80	50
Bea	50	80
<i>Equality under Risk</i>		
Ann	80	50
Bea	80	50
<i>Equality under Certainty</i>		
Ann	$65 - c$	$65 - c$
Bea	$65 - c$	$65 - c$

On our egalitarian view, this alternative is superior to the preceding two. For by ensuring that all are in the same boat, it eliminates all unfair inequality without loss of expected total well-being.

Finally, suppose that the following alternative also becomes available:

Equality under Certainty: This treatment will improve both Ann's and Bea's condition to that of a merely partial, but still substantial, visual impairment. We will consider both cases in which the level of well-being associated with this partial impairment is precisely halfway between the well-being associated with complete blindness and a full cure and cases in which this level falls short of this halfway point. The shortfall is given by a cost c , with $0 \leq c < 15$.

On our egalitarian view, if $c = 0$, then *Equality under Certainty* is of course better than the first two inequalitarian alternatives, because it eliminates inequality at no cost in expected total well-being. Moreover, *Equality under Certainty* is precisely as good as *Equality under Risk*. In the absence of inequality, this form of egalitarianism simply tells us to choose a best prospect for each individual; because, for $c = 0$, both *Equality under Risk*

and *Equality under Certainty* offer each individual an expected well-being of 65, both are equally good prospects for each person.

For a sufficiently small, positive cost ($c > 0$), *Equality under Certainty* will still be chosen over the first two alternatives, because it eliminates all inequality with only a small reduction in expected total well-being. However, our egalitarian view will then regard it as inferior to *Equality under Risk*, because the latter offers each individual better prospects while ensuring equality.

III. A CAUTIOUS CRITERION FOR DECISIONS UNDER UNCERTAINTY

We shall now explain how we propose to extend our pluralist egalitarian view to cases of uncertainty. Let us start with a simple, one-person case. Suppose that Ann will go wholly blind unless she is treated. You must choose between providing Ann with an established risky treatment, which, given the extensive evidence available, you confidently believe has a 0.5 chance of curing her and a 0.5 chance of having no effect on her, and providing her with a novel, maximally uncertain treatment, which will either lead to a full cure or else be entirely ineffective. There is no information available on the probabilities associated with these possible outcomes of the uncertain treatment. Moreover, you do not possess precise prior beliefs about the probability of its effectiveness. Which treatment(s) is it morally permissible for you to provide? And which would you choose?

There is evidence that many decision-makers' answer to the latter question would be the merely risky treatment. Part of this evidence is that in a wide range of experiments involving self-interested choices, a large share of decision-makers (typically: a majority) strictly prefer a prospect in which they gain if a fair coin comes up heads to the same gain on an event about which they know only that its probability may be anything in a range from 0 to 1.¹⁶ They thereby display what is known as "uncertainty (or ambiguity) aversion" on their own behalf. (Those who are indifferent between this risky and uncertain prospect are commonly described as "uncertainty [ambiguity] neutral"; those who strictly prefer the uncertain prospect are known as "uncertainty [ambiguity] seeking.") And although there are less data on choices that concern others'

16. See Stefan Trautmann and Gijs van de Kuilen, "Ambiguity Attitudes," in *The Wiley Blackwell Handbook of Judgment and Decision Making*, ed. Gideon Keren and George Wu (Chichester: Wiley, 2015), pp. 89–116, table 1.

interests only, uncertainty aversion appears to be just as prevalent in such decisions.¹⁷

The argument for holding that this common uncertainty-averse attitude is morally and rationally permissible proceeds in two stages. The first stage pertains to belief formation. In this situation, by hypothesis, the combination of your evidence and prior beliefs is compatible with a wide range of assignments of probabilities to particular outcomes of the novel treatment. You therefore lack sufficient basis for a unique assignment of precise probabilities to the possible outcomes of this treatment. Indeed, to make such an assignment would seem to be arbitrary in the sense that it runs ahead of the information you have and ignores other possible assignments that are no less consistent with your prior beliefs and evidence. Neither rationality nor morality requires the formation of beliefs that lack sufficient foundation in the evidence.¹⁸ You are therefore not required to adopt a single precise assignment of probabilities to each possible outcome of the novel treatment. Instead, it is reasonable for you to take account of the full range of probability assignments that are supported by the data and your prior beliefs. In other words, you may consider everything from the worst probability distribution over the outcomes “wholly ineffective” and “full cure” that is consistent with your evidence and prior beliefs (viz., that the novel treatment provides Ann with *no* chance of a cure) to the best probability distribution consistent with this information and these beliefs (viz., that it is *sure* to cure her), without reducing them to a single probability distribution.

The second stage pertains to how this range of probability assignments over pertinent outcomes can permissibly figure in your decision-making. The central idea is that although you should assign some decision weight to both the worse and better possible probability distributions over outcomes, how much decision weight to assign to each is, within a considerable range of sensible weights, up to you. Cautiously assigning somewhat greater decision weight to the worse possible probability distributions than to the better ones is in this

17. See Christian Koenig-Kersting and Stefan Trautmann, “Ambiguity Attitudes in Decisions for Others,” *Economics Letters* 146 (2016): 126–9 and Saposnik et al., “Decision-Making in Multiple Sclerosis.”

18. James M. Joyce, “How Probabilities Reflect Evidence,” *Philosophical Perspectives: Epistemology* 19 (2005): 153–78, at pp. 168–71; Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler, “Is It Always Rational to Satisfy Savage’s Axioms?” *Economics and Philosophy* 25 (2009): 285–96.

sensible range.¹⁹ Such caution in the face of an inability to arrive at precise probabilistic assignments is the central motivation for uncertainty aversion.

Our claim is not that uncertainty aversion is the *only* reasonable attitude. It is merely that a moderate degree of such aversion is perfectly sensible. Caution of the kind outlined could, we believe, be offered as a good reason for a choice of the risky treatment over the uncertain treatment to anyone concerned with Ann’s welfare.

We note, however, that despite its appeal, the rationality of uncertainty aversion is controversial among decision theorists. The reason is that if one assumes, as we have done, that under conditions of risk it is rationally required to obey the von Neumann–Morgenstern axioms, then uncertainty aversion and a central axiom of decision theory, the Sure Thing Principle, cannot be reconciled.²⁰ Moreover, violation of the Sure Thing Principle has unpalatable implications.²¹

19. See Ellsberg, “Risk, Ambiguity, and the Savage Axioms,” 643–69, and *Risk, Ambiguity, and Decision* (New York: Taylor & Francis, 2001); Binmore, *Rational Decisions*; Gilboa et al., “Is It Always Rational to Satisfy Savage’s Axioms?,” 285–96, and Hare, “Risk and Radical Uncertainty in HIV Research,” 87–9.

20. However, see Richard Bradley, *Decision Theory with a Human Face* (Cambridge: Cambridge University Press, 2017), sec. 9.5, for an argument that uncertainty aversion can be made consistent with the Sure Thing Principle if one drops the von Neumann–Morgenstern axioms.

21. The Sure Thing Principle is due to Leonard Savage, *The Foundations of Statistics*, 2nd ed. (New York: Dover, 1972). It can be understood as follows (see Bradley, *Decision Theory*, sec. 4.4). Consider Table n1, in which f , g , f' , and g' are alternative acts, E and *not- E* are events, and X , X^* , Y , and Y^* are constant outcomes (i.e., X is the same outcome no matter whether generated through a combination of choosing f and the occurrence of E or through choosing f' and the occurrence of E , etc. for X^* , Y , and Y^*).

Table n1.

Alternative	Event	
	E	<i>not-E</i>
f	X	Y
g	X^*	Y
f'	X	Y^*
g'	X^*	Y^*

Intuitively, f should be preferred to g just in case X should be preferred to X^* . This is because f and g have the same outcome whenever E is not the case, and therefore should be evaluated solely in terms of their outcomes when E is the case. Consequently, any other alternatives such as f' and g' , which have the same outcomes as f and g , respectively, whenever E is the case, and identical outcomes when it is not, should be ranked in the same order as f and g . In other words, the Sure-Thing Principle states that, in Table n1, f is preferred to g if and only if f' is preferred to g' . For an explanation of why uncertainty aversion can lead one

The debate on whether one should conclude that uncertainty aversion is irrational is extensive. Rather than reviewing this debate, we will simply state our take on it, which is that there is a tension between independently attractive principles of rationality, including, on the one hand, that rationality does not require a decision-maker to posit precise probabilities when they lack sufficient grounds for doing so and that a decision-maker is entitled to be cautious in the face of such a lack, and, on the other hand, that a decision-maker should respect other attractive principles of rational choice. Different ways of “trading off” such incompatible ideal standards of rationality are sensible, and among the sensible ways of making trade-offs are uncertainty-averse decision principles.²²

In the remainder of this article, we will therefore explore what would follow if a degree of uncertainty aversion were both rationally and morally acceptable. This question is worth exploring because uncertainty aversion strikes us, many everyday decision-makers, and a considerable number of experts as a reasonable attitude, and it gives rise to underexplored issues of distributive justice.

Many uncertainty-averse decision criteria have been proposed. For illustrative purposes, we will here use a simple but popular criterion first put forward by Leonard Hurwicz and later developed together with Kenneth Arrow, which pays attention to only the least favorable and most favorable probability distributions that are consistent with the decision-maker’s information and prior beliefs.²³ Our conclusions hold for all other leading criteria, including those that give some weight to all probability

to violate this principle and for discussion of the problems to which this violation gives rise, see Nabil Al-Najjar and Jonathan Weinstein, “The Ambiguity Aversion Literature: A Critical Assessment,” *Economics and Philosophy* 25 (2009): 249–84 and Fleurbaey, “Welfare Economics, Risk and Uncertainty,” *Canadian Journal of Economics* 51 (2018): 5–40, at pp. 29–32.

22. Here, we are agreeing with Gilboa et al., “Is It Always Rational to Satisfy Savage’s Axioms?”; Marciano Siniscalchi, “Two Out of Three Ain’t Bad: A Comment on ‘The Ambiguity Aversion Literature: A Critical Assessment’,” *Economics and Philosophy* 25 (2009): 335–56; Heal and Millner, “Uncertainty and Decision Making,” at pp. 129–30 and “Uncertainty and Ambiguity,” at pp. 23–4.

23. Leonard Hurwicz, “Optimality Criteria for Decision Making under Ignorance,” *Cowles Commission Discussion Paper, Statistics* 370 (1951), and Arrow and Hurwicz, “An Optimality Criterion for Decision Making under Ignorance,” in *Uncertainty and Expectations in Economics*, ed. Charles Carter and J. Ford (Oxford: Basil Blackwell, 1972), pp. 1–11. For discussion of this criterion, see Binmore, *Rational Decisions*; and Peter Wakker, *Prospect Theory for Risk and Uncertainty* (Cambridge: Cambridge University Press, 2010), sec. 11.5.

distributions that the decision-maker regards as consistent with their evidence and beliefs.²⁴

On what is known as the α -Hurwicz or α -maxmin criterion, one values each person's prospect at $\alpha \times$ its expected value given the least favorable probability distribution consistent with one's information and prior beliefs plus $(1 - \alpha) \times$ its expected value given the most favorable probability distribution that is so consistent, where $0 \leq \alpha \leq 1$ is the Hurwicz pessimism-optimism index. Uncertainty aversion involves giving more decision weight to the least favorable possible probability distribution than to the most favorable one; in other words, it involves taking $\alpha > 0.5$. (The criterion reduces to orthodox decision theory when a decision-maker uses a single probability distribution.) In what follows, we will assume a decision-maker who has a fixed, permissible degree of uncertainty aversion both when they evaluate a prospect for the sake of a single individual and when they evaluate a multi-person prospect. This implies a constant $\alpha > 0.5$ for all decisions.

By way of illustration, consider the experimental treatment with which we opened this section and which we represented by a case in which Ann is cured if and only if a red ball is drawn from a wholly uncertain urn. An uncertainty-averse decision-maker who uses the α -maxmin criterion will consider both the most pessimistic assessment of the information available—according to which there are no red balls in this urn—and the most optimistic assessment—according to which it contains only red balls. Moreover, they will give at least somewhat greater weight to the former than to the latter. Because of this cautious form of evaluation, they will regard the uncertain treatment as less good for Ann, in prospect, than giving her a risky treatment which would carry a 0.5 probability of a full cure and a 0.5 probability of being wholly ineffective. For example, a moderately uncertainty-averse decision-maker for whom $\alpha = 0.6$ will regard Ann's wholly uncertain prospect as equivalent to a treatment with an expected value of 62 units of well-being or 3 units of expected well-being less than this risky treatment. (Despite the fact that the criterion permits us to assign such equivalents to uncertain prospects, the value of an uncertain prospect when applying this criterion is not an expected value, because the decision

24. Our conclusions hold, for example, for the criteria proposed in Gilboa and Schmeidler, "Maxmin Expected Utility with Non-Unique Prior," and Klibanoff et al. "A Smooth Model of Decision Making under Ambiguity."

weights applied to different possible probability distributions are not probabilities. When we are discussing uncertain and/or risky prospects, we therefore use the more general term “prospective value.”)

Although in this simple case of Ann’s experimental treatment this criterion gives weight to both the worst and best possible *outcome*, this is only because, in this example, the most pessimistic assignment of probabilities is a certainty of failure and the most optimistic assignment of probabilities is a certainty of a cure. Whenever the decision-maker can rule out such extreme probability distributions, α -maxmin gives weight not to the worst and best *outcomes*, but to the lowest and highest *expected values* that the decision-maker assigns to the prospect. By way of illustration, suppose that the decision-maker gained more information about this novel treatment, so that the uncertainty involved was reduced as follows: they now rationally conclude the treatment has between a 0.25 and 0.75 chance of curing Ann. The α -maxmin criterion evaluates this revised prospect as follows: α times the most pessimistic assessment of its expected value (viz., that there is a 0.25 chance of curing Ann) plus $(1 - \alpha)$ times the most optimistic assessment of its expected value (that there is a 0.75 chance of curing Ann). Or, filling in the numbers:

$$\alpha(0.25 \times 80 + 0.75 \times 50) + (1 - \alpha)(0.75 \times 80 + 0.25 \times 50).$$

A moderately uncertainty-averse decision-maker for whom $\alpha = 0.6$ will therefore regard this partly uncertain treatment as equivalent to a treatment with an expected value of 63.5 units of well-being or precisely in between the value of the aforementioned wholly uncertain treatment and the value of the aforementioned merely risky treatment, which has a 0.5 chance of effecting a cure. This illustrates that, on this criterion, reducing the range of uncertainty also, naturally, reduces the depressing effect it has on the value of a prospect.

IV. WHEN REDUCING UNCERTAINTY AND LIMITING INEQUALITY ARE CONGRUENT

We will now review ways in which adding uncertainty aversion to our egalitarian view generates novel implications. We first focus on cases in which the aim of reducing uncertainty does not conflict with the aim of reducing inequality. (We deal with conflicts between these aims in the next section.)

Table 2.

Alternative	Event	
	Draw from a risky urn	
	red_r ($p_{red} = 0.5$)	$black_r$ ($p_{black} = 0.5$)
<i>Equal Risk, Unequal Final Well-being</i>		
Ann	80	50
Bea	50	80
Alternative	Draw from an uncertain urn	
	red_u ($0 \leq p_{red} \leq 1$)	$black_u$ ($0 \leq p_{black} \leq 1$)
	<i>Equal Uncertainty, Unequal Final Well-being</i>	
Ann	80	50
Bea	50	80

Suppose that a decision-maker must choose between the aforementioned *Equal Risk, Unequal Final Well-being* and the following:

Equal Uncertainty, Unequal Final Well-being: This treatment will either cure Ann and leave Bea wholly blind or cure Bea and leave Ann wholly blind, with no information available about the probability of either outcome.

These alternatives are displayed in Table 2. We use red_u ($black_u$) to signify the event of a red (black) ball being drawn from an uncertain urn.

Our pluralistic view requires that we take account of both the distribution of individuals' prospects and the prospective value of the possible anonymized distributions of final well-being. Let us consider each in turn. The risky alternative ensures equality of prospects, as does the uncertain alternative. However, the uncertain alternative gives each individual a less valuable prospect than its risky counterpart. Considering individuals' prospects, therefore, *Equal Risk, Unequal Final Well-being* is clearly superior. Furthermore, in terms of the prospective value of the possible distributions of final well-being, the two are equivalent. For, under each of these alternatives, the anonymized distribution of final well-being is certain: one person will be fully cured, another will go wholly blind. One can therefore say that although one of these alternatives contains *individual-level* uncertainty, neither contains

Table 3.

Alternative		Event	
		Draw from a risky urn	
		red_r ($p_{red} = 0.5$)	$black_r$ ($p_{black} = 0.5$)
<i>Equality under Risk</i>			
	Ann	80	50
	Bea	80	50
<i>Equality under Uncertainty</i>		Draw from an uncertain urn	
		red_u ($0 \leq p_{red} \leq 1$)	$black_u$ ($0 \leq p_{black} \leq 1$)
	Ann	80	50
	Bea	80	50

any *population*-level uncertainty. All things considered, *Equal Risk, Unequal Final Well-being* is therefore more choiceworthy, but only because it avoids the depressing effect of uncertainty on the value of individuals' prospects.

Now imagine a choice between the aforementioned *Equality under Risk* and the following:

Equality under Uncertainty: This treatment will either cure both children or leave them both to go wholly blind, with no information available about the probability of either outcome.

These alternatives are depicted in Table 3. In this case, both a concern for individuals' prospects and a concern for the prospective value of the possible anonymized distributions of final well-being point in the same direction. *Equality under Uncertainty* offers each individual a less valuable prospect. Moreover, it generates population-level uncertainty, because the decision-maker is uncertain about the anonymized distribution of final well-being. This lowers the value of *Equality under Uncertainty*, because the worst possible probability distribution (i.e., that the probability that both individuals are cured is 0) receives greater weight than the best possible probability distribution (i.e., that the probability that both are cured is 1). An uncertainty-averse view will therefore have two reasons for judging that it is better to opt for *Equality under Risk*.

Table 4.

Alternative	Event			
	Draw from wholly uncertain urn for Ann, risky urn for Bea			
	red_w, red_r	$red_w, black_r$	$black_w, red_r$	$black_w, black_r$
<i>Unequal Uncertainty</i>				
Worst prob. distrib. ²⁵	0	0	0.5	0.5
Best prob. distrib.	0.5	0.5	0	0
Ann	80	80	50	50
Bea	80	50	80	50
Draws from distinct, partly uncertain urns for Ann and Bea				
	red_w, red_u	$red_w, black_u$	$black_w, red_u$	$black_w, black_u$
<i>Equal Moderate Uncertainty</i>				
Worst prob. distrib. ²⁶	0.0625	0.1875	0.1875	0.5625
Best prob. distrib.	0.5625	0.1875	0.1875	0.0625
Ann	80	80	50	50
Bea	80	50	80	50

So far, we have analyzed cases in which, while keeping inequality constant, a decision-maker can ensure less uncertainty. Now, we will consider a case in which, keeping total uncertainty constant, a decision-maker can equalize its burden.

Suppose that a decision-maker must choose between the following:

Unequal Uncertainty: Ann is given a novel treatment which will either cure her or instead leave her wholly blind, with no information about the probability of either outcome. Bea is given a distinct treatment which will either, with probability 0.5, cure her or, with probability 0.5, leave her wholly blind.

Equal Moderate Uncertainty: Ann and Bea are each given different distinct, moderately uncertain treatments, each of which will either offer a

25. One arrives at the worst (best) distribution by assuming the smallest (largest) share of red balls in Ann's uncertain urn consistent with the available evidence, namely, 0 (1).

26. One arrives at the worst (best) distribution by assuming the smallest (largest) share of red balls in Ann's and Bea's uncertain urns consistent with the available evidence, namely, 0.25 (0.75) in each urn.

full cure or leave its recipient wholly blind. For each of their treatments, the probability of a cure ranges from 0.25 to 0.75.

These alternatives are represented in Table 4.

The choice between *Unequal Uncertainty* and *Equal Moderate Uncertainty* can be thought of as follows. Ann and Bea each face a draw from a separate urn. Each receives a cure if a red ball is drawn from their urn; if a black ball is drawn, their treatment is ineffective. The decision-maker must fill each urn with 100 balls. They have four bags of 50 balls each: two risky bags containing an equal mix of red and black balls and two wholly uncertain bags about which the decision-maker has no information except that they can be any proportion of red and black. If they empty both uncertain bags into Ann's urn and both risky bags into Bea's urn, then they generate *Unequal Uncertainty*. By contrast, if they fill each urn with one uncertain and one risky bag, then they generate *Equal Moderate Uncertainty*. The former places all the burden of uncertainty on Ann's prospects. By contrast, the latter equalizes the burden of uncertainty. Moreover, it is natural to suppose that the total burden created by the uncertain balls is not increased when they are divided equally.²⁷ From the perspective of the distribution of the value of individuals' prospects, therefore, *Equal Moderate Uncertainty* is clearly superior.

We must also consider the prospective value of the possible anonymized distributions of final well-being associated with each of these alternatives. Using the α -maxmin criterion, it is sufficient to consider only the worst and best among the possible probability distributions, which are listed in Table 4. Now, under *Unequal Uncertainty*, in both the pessimistic scenario (i.e., Ann's urn contains no red balls) and the optimistic scenario (i.e., Ann's urn contains red balls only), an unequal outcome—in which only one person is cured—has a probability of 0.5. (This is the sum of the probabilities for the events $\{red_w, black_r\}$ and $\{black_w, red_r\}$.) By contrast, under *Equal Moderate Uncertainty*, in both the pessimistic and optimistic scenario, the probability of an unequal

27. Indeed, according to the α -maxmin criterion, under *Unequal Uncertainty*, the value of Ann's prospects is $50\alpha + 80(1 - \alpha)$ and the value of Bea's prospects is 65. Under *Equal Moderate Uncertainty*, the value of Ann's prospects is $\alpha(0.25 \times 80 + 0.75 \times 50) + (1 - \alpha)(0.75 \times 80 + 0.25 \times 50)$; the same is true of Bea's prospects. In both cases, the total value is therefore $145 - 30\alpha$. An egalitarian view that incorporates α -maxmin will therefore hold *Equal Moderate Uncertainty* is superior because it distributes this total more equally.

Table 5.

Alternative	Event	
	Draw from an uncertain urn	
	red_u ($0 \leq p_{red} \leq 1$)	$black_u$ ($0 \leq p_{black} \leq 1$)
<i>Equal Uncertainty, Unequal Final Well-being</i>		
Ann	80	50
Bea	50	80
<i>Equality under Certainty</i>		
Ann	$65 - c$	$65 - c$
Bea	$65 - c$	$65 - c$

outcome is 0.375. In other words, *Equal Moderate Uncertainty* makes an unequal outcome less likely no matter whether the odds are stacked against the person(s) facing uncertainty or whether the odds are in their favor. This makes it better from the perspective of the prospective value of the anonymized distribution of final well-being.²⁸ We can conclude that our uncertainty-averse egalitarian view yields the plausible verdict that one should distribute the burden of uncertainty equally.

Our view does not merely posit a novel object of egalitarian concern (the disvalue of uncertainty); it also lends additional force to the egalitarian aim of directing aid toward those who end up less well off than others. By way of illustration, suppose that our decision-maker must choose between *Equal Uncertainty*, *Unequal Final Well-being* and *Equality under*

28. More precisely, and writing $v\{80, 50\}$ for the value of a distribution in which one person gets 80 and another one gets 50, the prospective value that the α -maxmin criterion assigns to the distribution of final well-being under *Equal Moderate Uncertainty* minus the prospective value that it assigns to the distribution of final well-being under *Unequal Uncertainty* is:

$$0.0675 \times (v\{80, 80\} - v\{80, 50\}) - 0.0675 \times (v\{80, 50\} - v\{50, 50\}).$$

This expression will be positive just in case “leveling up” from (80, 50) to (80, 80) generates more moral value than “leveling down” from (80, 50) to (50, 50) destroys. And, this will be true on our pluralistic egalitarian view. This is because the added 30 units of well-being in the leveling up scenario are valuable because they are good for a person and reduce inequality. By contrast, the disvalue of the lost 30 units of well-being in the leveling down scenario is tempered by the fact that their loss reduces inequality.

Certainty. For convenience, both are represented in Table 5. Recall that c is the cost of achieving both equality and certainty, with $0 \leq c < 15$.

Let us again consider both the value of individuals' prospects and the prospective value of the anonymized distribution of final well-being. Under *Equal Uncertainty*, *Unequal Final Well-being*, the value of each individual's prospect is depressed by the fact that the decision-maker has no information about their chance of ending up disadvantaged. The badness of this uncertainty for an individual is determined by the gap between the expected value of this individual's prospects given the possible probability distribution that is least favorable to them and the expected value of this individual's prospects given the possible probability distribution that is most favorable to them. *Equality under Certainty* is valuable because it altogether eliminates this gap. From the perspective of the value of individuals' prospects, an uncertainty-averse decision-maker should therefore be willing to incur a cost ($c > 0$) to eliminate this uncertainty.

Turning to the prospective value of the possible distributions of final well-being, a drawback of *Equal Uncertainty*, *Unequal Final Well-being* is, naturally, the certainty of outcome inequality. An inequality-averse decision-maker will therefore be willing to pay a price ($c > 0$) to eliminate this inequality.

Overall, both uncertainty aversion and inequality aversion will prompt us to incur a cost to remove inequality. Moreover, jointly, they will justify paying a higher price to achieve equality than either alone would. To see why, suppose for the moment that our decision-maker remained inequality averse but became indifferent to uncertainty (i.e., their $\alpha = 0.5$). They would then evaluate *Equal Uncertainty*, *Unequal Final Well-being* as equivalent to *Equal Risk*, *Unequal Final Well-being* (the latter is described in Table 1). Suppose that to eliminate the inequality of final well-being in these alternatives, it is right to incur up to, but no more than, a cost to each person of c^* units of expected well-being. We can then say that, for an uncertainty-neutral, but inequality-averse decision-maker, both *Equal Uncertainty*, *Unequal Final Well-being* and *Equal Risk*, *Unequal Final Well-being* are equivalent to giving each of Ann and Bea $65 - c^*$ for certain.

Now suppose that our decision-maker regained their uncertainty aversion. They would then regard *Equal Uncertainty*, *Unequal Final Well-being* as worse than *Equal Risk*, *Unequal Final Well-being*. By transitivity, they would regard the uncertain alternative as worse than giving both Ann and Bea $65 - c^*$ for certain. In other words, they would regard *Equal*

Uncertainty, Unequal Final Well-being as equivalent to *Equality under Certainty* only for a cost larger than c^* . It follows that an uncertainty-averse egalitarian view justifies incurring a larger cost to achieve both equality and certainty than an uncertainty-neutral egalitarian view would countenance.

Let us summarize the distinctive implications of our view uncovered in this section. First, and straightforwardly, it will favor situations in which a better basis is available for assigning probabilities to outcomes. This is illustrated by the stylized choices in Tables 2 and 3. In real-world cases, the view will therefore display a tendency to favor policies with an extensive evidence base over ones with a minimal evidence base, keeping other things equal. Under these circumstances, it will also favor familiar over unfamiliar treatments for patients and make the provision of the latter harder to justify.²⁹

Second, the proposed view posits an additional object of egalitarian concern, namely, the burden of uncertainty. In general, it implies that there is unfairness in situations in which some face a greater burden of uncertainty than others, either because there is less information about the likelihood of the possible threats they face or because there is much more at stake for them. The case outlined in Table 4 provides a stylized example. A realistic case in which such inequality is of concern is climate policy. For it is likely that people in poorer nations who inhabit marginal lands and who are dependent on the weather for their livelihood face larger burdens of uncertainty than urbanites in wealthy countries.

Third, in uncertain situations in which we know that one person's good fortune will be the counterpart of another's misfortune, uncertainty aversion and inequality aversion point us in the same direction. In such situations, steering benefits to whoever turns out to be less fortunate reduces the stakes for each person and thereby lessens the burden of uncertainty; of course, it also reduces inequality in final well-being. Table 5 provides an example where this process of improving the lot of the worst off can be pursued to the point of equality, but the impulse toward equality will be present even when perfect equality cannot be reached. A policy issue to which this may be relevant is the taxation of returns from complex financial instruments held by pension funds on behalf of ordinary workers. These instruments are uncertain prospects that will yield both winners and losers. Measures that dampen the variability of returns (e.g., taxing gains and allowing a tax

29. Cf. Saposnik et al., "Decision-Making in Multiple Sclerosis."

Table 6.

Alternative	Event	
	Draw from an uncertain urn	
	red_u ($0 \leq p_{red} \leq 1$)	$black_u$ ($0 \leq p_{black} \leq 1$)
<i>Equal Uncertainty, Unequal Final Well-being</i>		
Ann	80	50
Bea	50	80
<i>Equality under Uncertainty</i>		
Ann	80	50
Bea	80	50

deduction for losses) will be valuable both because they reduce uncertainty and because they decrease inequality in final well-being.³⁰

V. WHEN REDUCING UNCERTAINTY AND LIMITING INEQUALITY ARE AT ODDS

We now turn to cases in which uncertainty aversion and inequality aversion pull in opposite directions. By way of illustration, imagine a choice between the aforementioned alternatives *Equal Uncertainty, Unequal Final Well-being* and *Equality under Uncertainty* as depicted in Table 6.

In this case, in terms of the prospects they grant individuals, both alternatives are equivalent. However, in terms of the prospective value of the possible distributions of final well-being they generate, matters are less straightforward. If the decision-maker chooses *Equal Uncertainty, Unequal Final Well-being*, they can be certain that one person will be cured and one will go wholly blind. But this absence of population-level uncertainty about the anonymized distribution of final well-being comes at the cost of guaranteed inequality. If, instead, they choose *Equality under Uncertainty*,

30. Although some investors in such instruments may be adequately informed and wholly responsible for their choice of investments (so that a responsibility-sensitive egalitarian may have attenuated reason for being concerned with inequalities that result), many of the returns also come to participants in funds who should not be held fully responsible for the relevant investment decisions (see, e.g., Michael Lewis, *The Big Short* [New York: Norton, 2010], chap. 6, recounting how, among others, Japanese farmers' unions and European pension funds were exposed to swings in the value of collateralized debt obligations). Inequalities in final well-being between such investors should concern egalitarians.

Table 7.

Alternative	Event	
	Draw from an uncertain urn	
	red_u ($0 \leq p_{red} \leq 1$)	$black_u$ ($0 \leq p_{black} \leq 1$)
<i>Equality under Uncertainty</i>		
Ann	80	50
Bea	80	50
<i>Equality under Certainty</i>		
Ann	$65 - c$	$65 - c$
Bea	$65 - c$	$65 - c$

there will be guaranteed equality, because Ann and Bea will either both go wholly blind or both be cured. But the fact that they will either sink or swim together generates substantial uncertainty regarding the value of the distribution of final well-being.

The view we are developing is silent on the choice between these alternatives. It assumes only that some degree of inequality aversion is required and that some degree of uncertainty aversion is both permissible and adopted by the decision-maker. Consistently with these commitments, decision-makers may decide this case differently. A decision-maker who is strongly inequality averse but only slightly uncertainty averse will choose *Equality under Uncertainty*; a decision-maker who is only slightly inequality averse but very averse to uncertainty will favor *Equal Uncertainty*, *Unequal Final Well-being*. Here, we will not argue that one of these ways of trading off the two concerns is uniquely right; the key conclusion is simply that our view incorporates some resistance to the egalitarian impulse to bind people’s fates together.

One further important lesson about how uncertainty aversion changes egalitarianism is this. Under risk, our egalitarian view embodies a tendency to allocate benefits away from the fortunate and toward the unfortunate *when and only when the former are (or would be) better off than others and the latter are (or would be) worse off than others*. This is illustrated by the fact that in choosing between *Equality under Risk* and *Equality under Certainty* in Table 1, our view is unwilling to pay any cost to redistribute from the fortunate potential futures of Ann and Bea to their less fortunate potential futures. By contrast, under uncertainty, our view embodies a

tendency to direct benefits from the fortunate toward the unfortunate 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*, which are both depicted in Table 7. Due to uncertainty aversion, for some sufficiently small, positive c , *Equality under Certainty* is superior to *Equality under Uncertainty*; this is indicative of the cost our view is willing to pay to reduce the burden of uncertainty by directing benefits away from Ann and Bea's better potential futures toward their worse potential futures.

To be sure, on our view, the strength of this tendency to allocate benefits away from the fortunate and toward the unfortunate will be strongest when it both reduces inequality between people and reduces the disvalue of uncertainty for each person. Nonetheless, the predilection in uncertain situations to direct resources toward all victims of misfortune is striking. It follows from this view that we have special reason to make provision for possible adversity to which we cannot assign a precise probability, even if this adversity would, if it occurred, equally affect everyone in the collectivity with which we are concerned. A realistic case may be the purchase of a financial hedge for uncertain and volatile revenues from a country's collectively owned natural resources, to protect against a downturn that would depress all citizens' livelihoods.³¹

VI. UNCERTAINTY, PARETO, AND ANTI-EGALITARIANISM

We shall now demonstrate that the permissibility of uncertainty aversion should lead us to reassess a common argument against the kind of egalitarianism we have been developing. Consider the following principle:

Pareto for Prospects: If, for every person, a first alternative provides a more valuable prospect than a second, then the first alternative should be chosen over the second.

31. If, the quotation from Keynes in Section I suggests, one were to regard the oil price as uncertain, then the Mexican government's hedging program may provide a concrete example. See Reuters News Agency, "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 August 12, 2017.

The motivation for this principle is that a social decision-maker should choose an alternative that could be chosen for the sake of each person, if such an alternative exists. The principle holds that an alternative is more choiceworthy for a person's sake just in case it has greater prospective value for that person.

It is well known that the egalitarian view under consideration can violate Pareto for Prospects. By way of illustration, consider a choice between *Equality under Certainty* and *Equal Risk, Unequal Final Well-being* (see Table 1). For a sufficiently small cost c , our pluralist egalitarian view chooses the former, because it eliminates all forms of unfair inequality, although it yields less valuable prospects for each person. Some regard this as a reason to favor utilitarianism, which, they allege, does not conflict with Pareto for Prospects.³²

However, if, as we have argued, uncertainty aversion is permissible, then utilitarians are not well placed to advance this argument, because utilitarianism will also violate Pareto for Prospects. Consider the choice between *Equal Uncertainty, Unequal Final Well-being* and *Equality under Certainty* outlined in Table 5. An uncertainty-averse evaluation of each individual's prospects yields the result that for some sufficiently small, positive c , *Equality under Certainty* gives each individual more valuable prospects. Pareto for Prospects therefore holds that for a sufficiently small, positive c , one must choose *Equality under Certainty*. However, for all $c > 0$, *Equal Uncertainty, Unequal Final Well-being* is guaranteed to yield more total well-being, so that a utilitarian must choose it.³³

32. See, for example, Harsanyi, "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy* 63 (1955): 309–21, and McCarthy, "Distributive Equality." For a response to this objection to egalitarianism, see Fleurbaey and Voorhoeve, "Decide as You Would with Full Information! An Argument against ex ante Pareto," in *Inequalities in Health: Concepts, Measures and Ethics*, ed. Nir Eyal, Samia Hurst, and Ole Frithjof Norheim (Oxford: Oxford University Press), pp. 113–28.

33. Indeed, assuming uncertainty aversion, this case demonstrates a tension between Pareto for Prospects and *all* views about the value of an anonymized distribution of final well-being. To see why, assume a specific degree of uncertainty aversion by the decision-maker on an individual's behalf. This degree will imply a specific, positive cost for which it is, prospectively, just as good for an individual to face *Equality under Certainty* as it is to face *Equal Uncertainty, Unequal Final Well-being*. Call this cost c^* (with $c^* > 0$) (e.g., using the α -maxmin criterion, if $\alpha = 0.6$, then $c^* = 3$). For any cost larger than c^* , Pareto for Prospects will demand the choice of *Equal Uncertainty, Unequal Final Well-being*. For any cost smaller than c^* , Pareto for Prospects will demand the choice of *Equality under Certainty*. Now consider the value of the anonymized distribution of well-being associated with these alternatives. For both alternatives, this distribution is certain. The question is therefore how the

We draw two conclusions from these findings. The first concerns the dialectic between utilitarians and egalitarians: only those utilitarians who deny the permissibility of uncertainty aversion can consistently reject egalitarianism simply because it violates Pareto for Prospects. Utilitarians who grant the permissibility of uncertainty aversion might, of course, offer the following, weaker objection to egalitarianism. "Egalitarianism violates Pareto for Prospects under conditions of both risk and uncertainty, whereas utilitarianism violates it under conditions of uncertainty alone. It counts in favor of a view if it violates it in a more limited set of conditions. Therefore, utilitarianism is, in one respect at least, superior to egalitarianism." This objection, however, grants that violating Pareto for Prospects is not a deal breaker. And once one is permitted to violate this principle for the sake of improving total final well-being, why would one not also be permitted to do so for the sake of reducing inequality?

Our second conclusion is substantive. By focusing only on each individual's prospects, considered separately, Pareto for Prospects is insensitive to how combinations of these prospects give rise to possible patterns of final well-being. But these patterns matter, both because they determine the fairness of the eventual outcome and because they determine how much uncertainty about the value of the distribution of final well-being a social decision-maker faces. After all, a decision-maker is in a very different situation when some will definitely end up better off than others than when no such inequality can arise. They are also in a different situation when they can accurately predict the impact of their policies at a population level than when they must contend with a wide range of possible population-level outcomes and are not in a position to assign precise probabilities to these outcomes. That Pareto for Prospects does not

value of a distribution in which one person has 80 and another one has 50 compares with the value of a distribution in which both people have $65 - c$. Every theory about the value of a distribution of final well-being will have a cost for which these two are equivalent. For a utilitarian, as we have seen, this is $c = 0$; for a pluralist egalitarian, it is some $c > 0$, with the size of the cost determined by the degree of inequality aversion; and so on for every other view. Now, whatever this distributive theory is, there is no reason that the cost at which it regards these two distributions of final well-being as equivalent should be precisely the cost c^* at which the two distributions offer each individual equivalent prospects. It follows that, if uncertainty aversion is permissible, any view that gives substantial weight to the value of anonymized distributions of final well-being will violate Pareto for Prospects. (We thank Richard Bradley for drawing our attention to this implication.)

permit decision-makers to take account of these differences is, we submit, a reason to reject it.

VII. CONCLUSION

Uncertain situations are ubiquitous. A common and, we have argued, permissible attitude in such situations is for a decision-maker to keep in mind the full range of distributions of probabilities over outcomes consistent with their limited evidence and prior beliefs and to respond cautiously to this range by giving some additional decision weight to the less favorable probability distributions within it. We have here explored what would follow if such an uncertainty-averse attitude was incorporated into a plausible egalitarian theory.

We highlight four implications of this view. First, and most straightforwardly, it favors risky alternatives over alternatives for which the decision-maker is not in a position to assign precise probabilities to possible outcomes. It therefore regards courses of action that avoidably involve uncertain possibilities of success and failure as difficult to justify. For example, it holds that it is more difficult to justify enrolling people in an experimental treatment when an alternative is available with well-established risks.³⁴

Second, our view identifies the disvalue of uncertainty as a new object of egalitarian concern. It can thereby justify special efforts to improve the prospects of those who face a larger burden of uncertainty, by improving information-gathering and making provision to improve their lot should the current, imprecisely estimated odds prove to be against them. A pertinent case is the possible effects of climate change on people whose lives and livelihoods would be threatened by large temperature rises or by changes in precipitation.

Third, in uncertain situations, uncertainty aversion serves as a counterweight to the egalitarian impulse to bind people's fates together. In such situations, ensuring that either everyone or no one is benefitted creates an uncertain possibility of collective misfortune, which the view will regard as especially problematic.

Fourth, in uncertain situations, the proposed view justifies directing benefits away from the fortunate and toward the unfortunate both when

34. Hare, "Risk and Radical Uncertainty in HIV Research."

they are different people and when they are merely different possible futures of the same person. It therefore adds force to the egalitarian impulse to resolve *interpersonal* trade-offs to the benefit of the less well-off. It also introduces a novel reason, namely, the reduction of the depressing effect of uncertainty on the value of an individual's prospects, to resolve *intrapersonal* trade-offs in a manner that makes special provision for the least fortunate potential future of the person concerned. This implication is relevant to justifications for a social safety net. As egalitarians have commonly understood this institution, it serves both to reduce inequalities and improve people's prospects by mitigating risks pertaining to income and health.³⁵ But where such a safety net must operate not merely under conditions of risk but also under uncertainty, our view provides the following additional reasons for its maintenance and expansion. By improving the fate of the worst off, a social safety net reduces both the depressing effect that individual-level uncertainty has on the value of individuals' prospects and population-level uncertainty about the value of the distribution of final well-being.

Finally, let us revisit our initial case involving the UK's policy in the face of the 2009 H1N1 pandemic. Analyses of the actual decision-making process reveal that key decision-makers in the civil service and the UK Government evaluated all courses of action exclusively in terms of their impact in the event of a devastating epidemic that would, if the more extensive of the contemplated policies were not implemented, claim between 65,000 and 750,000 lives.³⁶ Consequently, the government chose to invest the maximal amount in response to the threat. As it turned out, the virus was not particularly dangerous and the death toll was 457. Of course, this outcome alone does not show that the decision-making process at the time was flawed. However, some commentators have derided the UK Government's exclusive focus on the worst range of outcomes as excessively cautious and wasteful. They have suggested that, instead, it could have

35. See, for example, Nicholas Barr, *The Economics of the Welfare State*, 5th ed. (Oxford: Oxford University Press, 2012); Michael Otsuka, "How to Guard against the Risk of Living too Long: The Case for Collective Pensions," in *Oxford Studies in Political Philosophy*, Vol. III, ed. David Sobel, Peter Vallentyne, and Steven Wall (Oxford: Oxford University Press, 2017), pp. 229–51; and Alex Voorhoeve, "May a Government Mandate more Comprehensive Health Insurance than Individuals Want for Themselves?" in *Oxford Studies in Political Philosophy*, Vol. IV, ed. Sobel, Vallentyne and Wall (Oxford: Oxford University Press, 2018), pp. 167–91.

36. Hine, *An Independent Review of the UK Response to the 2009 Influenza Pandemic*; Oliver, "Ambiguity Aversion and the UK Government's Response to Swine Flu," 16–31.

made an “informed guess” assignment of precise probabilities to each possible outcome and then conducted a standard expected value calculation for each policy alternative.³⁷ This orthodox approach, however, has its own problems, because it demands that public decision-makers come up with a precise probability for every relevant outcome when even experts are unable to do so. Using such unfounded probabilities would seem hard to justify publicly. And in the absence of such a unique probability distribution, the proposed expected value analysis cannot be carried out.

The approach we have outlined here is, we submit, an attractive alternative to both the government’s exclusive focus on the worst case scenario and such orthodox expected value analysis. There are two points at which our proposed approach differs from “assuming the worst.” First, it focuses special attention on the least favorable probability distribution over outcomes consistent with available data and reasonable prior beliefs. That is, it focuses on the worst *expected value* rather than on the worst *outcome*. It is only in a case of complete ignorance that the worst expected value and the worst outcome are identical. As commentators have pointed out, the historical record of flu pandemics, alongside incoming information about the nature of the H1N1 virus, ensured the situation was not one of complete ignorance. This, they claim, made it unreasonable to reason as if the worst case was *certain* to occur.³⁸ If these observers are correct, then among all the probability distributions over outcomes in a sensible range, even the most negative would assign some probability to scenarios other than the worst.

Second, uncertainty aversion does not require giving decision weight only to the most negatively assessed expected value; it merely involves cautiously giving such gloomy estimates of expected value somewhat greater decision weight than more favorable assessments. Although this leaves a degree of freedom about these decision weights, a choice of such weights for high-stakes public decision-making under uncertainty strikes us as an unavoidable value judgment that is best openly debated.

37. See, for example, Oliver, “Ambiguity Aversion and the UK Government’s Response to Swine Flu,” 29–30.

38. See, for example, Oliver, “Ambiguity Aversion and the UK Government’s Response to Swine Flu,” 29–30; Hine, *An Independent Review of the UK Response to the 2009 Influenza Pandemic*, chap. 3.

Our critique of the UK Government's exclusive focus on the worst case does not, of course, show that substantial investment in pandemic preparedness was unwarranted. Indeed, our uncertainty-averse, egalitarian view identifies several reasons to consider such investment as valuable at the moment of decision. First, by improving individuals' outcomes in the event that the pandemic had turned out to be severe, it mitigated the depressing effect of uncertainty on the value of individuals' prospects. Second, limiting the number of deaths in this unfortunate eventuality alleviated population-level uncertainty. Finally, given that the impact of a severe epidemic on people's lives would have been highly disparate (with some, including many young people, dying and many more surviving), investment in saving lives if the pandemic had proven to be severe lowered the prospective degree of inequality in final well-being. Overall, besides demonstrating the value of an approach that treats decisions under uncertainty differently from merely risky decisions, this case highlights a central conclusion of our analysis, which is that uncertainty aversion both reinforces egalitarian reasons and provides new reasons to improve the lot of the unfortunate.