

The Moral & Rational Costs of Uncertainty Aversion

Brian Jabarian¹
May 2019

Abstract

In their article, “Egalitarianism under Severe Uncertainty”, (Philosophy and Public Affairs, 2018), Thomas Rowe and Alex Voorhoeve develop elegantly a theory of distributive justice, called “pluralist egalitarianism”, for cases under maximal uncertainty. In this précis, I firstly sketch their view and arguments. I then discuss their main cases. Finally, I suggest some points that we could discuss together in this very exciting 48-hour philosophy hackathon. Admittedly, these points won’t be exhaustive; you are welcome to suggest others in the comments.

§1 Rowe and Voorhoeve’s view

1.1 *Pluralist Egalitarianism*

Pluralist egalitarianism holds that

“[one] should improve people’s prospects for well-being, raise total well-being, and reduce inequality in both people’s prospects and in their final well-being (how well their lives end up going)” (Rowe & Voorhoeve, 2018, 243-244).

While one can find a comprehensive defence of pluralist egalitarianism elsewhere², the authors extend here this theory from cases under risk to cases under uncertainty. They build this extension by showing that it is a morally and rationally permissible distributive theory of justice under uncertainty.

Before discussing the challenges this extended theory faces, it is worth providing some background. Pluralist egalitarianism is built on two kinds of principles: moral and rational principles. The former principles are egalitarian and sketched in the definition above. The latter ones, however, need some elaboration. For cases under risk, these

¹jabarian@princeton.edu I thank Franz Dietrich, Marc Fleurbaey, Thomas Rowe and Alex Voorhoeve for their comments. I also thank Kévin Bertrand, Déborah Chekroun, Simon Lazarus, Anders Priergaard, Jimi Vaubien and the audience from the Social Ethics Discussion Seminar for useful discussions and reactions. I also gratefully thank the editors of PEA Soup, David Faraci and Peter M. Jaworski for setting up such a stimulating discussion.

²See Fleurbaey (2018), Fleurbaey and Voorhoeve (2018) and Otuska and Voorhoeve (2018).

rational principles are exclusively derived from standard decision theory. One cornerstone orthodox principle under risk relevant here is that the decision-maker is capable of making up their mind to form (or has access to) precise probabilities regarding the possible states of nature. For cases under uncertainty, the decision-maker loses this ability. Depending on the degree of severity of uncertainty, one ends up at first best, with a reasonable range of probabilities (cases of "moderate uncertainty") and at second best, with extreme intervals of probabilities (cases of "severe uncertainty"). At worst, it could result with no probability at all (cases of "maximal uncertainty"³). In these uncertain cases, an additional principle is required to guarantee that a decision is rational.

Several candidates are in the running to meet this rationality requirement in our context of distributive justice. The challenge is then to choose one that fulfils the two following conditions successfully. Firstly, the new principle has to be not only rationally permissible under uncertainty but also morally permissible (from a pluralistic egalitarian perspective). Secondly, this principle has to be flexible enough to accommodate a broader and dynamic framework of distributive justice to allow differential and (morally and rationally) permissible attitudes towards risk and uncertainty.

The new principle Rowe and Voorhoeve rely on is the uncertainty aversion principle. This latter holds that when choosing between a risky prospect and an uncertain prospect, one opts for the risky prospect⁴. Let's consider whether this principle is morally and rationally permissible. When considering its rational permissibility, one has to consider two dimensions: descriptive and normative. We can attest that the uncertain aversion principle is descriptively accurate since the empirical results show that this principle describes adequately how subjects behave uncertainty.

However, the normative rational permissibility of this uncertainty aversion attitude is still contested in the philosophical and economic theory literature. Namely, the debate has not yet reached a consensus whether a rational agent may permissibly display an uncertainty averse attitude. Rowe and Voorhoeve do not engage in this controversy. They instead adopt an assumption defended by some leading decision theorists according to which it is rationally permissible to display an uncertainty averse attitude, though it is not rationally required. If this debate has focused in length on the normative rational permissibility of this attitude, much less has been said about its moral permissibility. Rowe and Voorhoeve's paper is important because it fills this gap. And, it is original in developing a specific egalitarian interpretation of the uncertainty aversion.

To justify the moral permissibility of uncertainty aversion, the authors proceed as follows. They first propose a different meaning for "equality" than the standard one. Usually, equality is understood in terms of the outcome's values: two situations are equal if and only if individuals end up equally well off. One could extend this claim under uncertainty such that the only morally relevant information to define "equality" is still the outcome's

³Also known as "ambivalence" in the literature, see Bradley (2016).

⁴As initially shown in Ellsberg's example, see Ellsberg (1961).

values. However, according to the authors, this definition of equality, pertaining exclusively to value of final well-being, would rule out crucial moral information to design a fair system of distributive justice. Accordingly, one should incorporate the experience of uncertainty itself in the definition of equality under uncertainty. This integration can be seen as a moral benefit or a cost in the system of distributive justice.

For Rowe and Voorhoeve, facing uncertainty is a "burden" (op. cit. p. 242) in the sense of depressing the value of an individual's prospects. Therefore, it should correspond to a moral cost. Let us see why in the following situation. Suppose Ann will go wholly blind unless she is treated. As her doctor, you have two alternative treatments. The first treatment is well-known to all and risky. It has a 50% chance of curing her and 50% chance of having no effect on her. Since you have implemented it in the past, and so, have access to a small distribution of success and failure, you have certain prior beliefs in these current objective estimations of success and failure. The second treatment is entirely new and maximally uncertain. It leads to a full cure or no cure at all with no objective estimation of failure and success accessible. Since it is very new, you are not familiar enough to counter the absence of objective estimations by forming precise prior beliefs about its effectiveness. Despite leading to the same two possible levels of final well-being as the risky treatment, we can say that, in prospect, the uncertain treatment bears a moral cost, which, granting uncertainty aversion, would be morally impermissible to choose to incur on Ann's behalf, for anyone concerned by her welfare.

In sum, pluralistic, uncertainty-averse egalitarianism favours alternatives for which more fine-grained probabilistic information related to the states of nature is available. Besides, this view considers uncertainty as important moral information to rely on to take a fair distributive justice based decision and should count as a moral cost (in the sense of depressing the value of individuals' prospects) in the system of distributive justice.

1.2 Discussion of Cases

Let us now see how pluralistic egalitarianism works in a more concrete situation. Suppose Ann and Bea will go wholly blind and with a lifetime well-being of 50 (moderately good quality of life) unless they are treated. If fully cured, each individual would have a lifetime well-being of 80 (very good quality of life). Both are strangers to the decision-maker and each other. Unfortunately, the resources at the decision-maker's disposal do not suffice to cure fully both Ann and Bea. As listed below three classes of treatments are available to the decision-maker: certain, risky and uncertain treatments. The classes of certain and risky treatments each contain two different treatments. The class of uncertain treatments contains four different treatments. Among these eight treatments, the decision-maker has to choose one. This is where a distributive theory of justice can help to make up the decision-maker's mind.

- (1) *Inequality under Certainty*. Cure Ann and leave Bea to go wholly blind.
- (2) *Equal Risk, Unequal Final Well-being*. This treatment either cures Ann and is entirely ineffective for Bea (leaving her to go wholly blind) or, instead, is entirely ineffective for Ann (leaving her to go wholly blind) and cures Bea. These results are equally likely.
- (3) *Equality under Risk*. This treatment either cures both individuals or is wholly ineffective for both, with each result being equally likely.
- (4) *Equality under Certainty*. This treatment improves 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 $0 < c < 15$.
- (5) *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.
- (6) *Equality under Uncertainty*. This treatment either cures both individuals or leaves them both to go wholly blind, with no information available about the probability of either outcome.
- (7) *Unequal Uncertainty*. Ann is given a novel treatment which either cures her or instead leaves 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.
- (8) *Equal Moderate Uncertainty*. Ann and Bea are each given different distinct, moderately uncertain treatments, each of which will either offer a 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.

Then, the authors consider the following pairwise comparison of alternatives. They deduce the preference relations between the alternatives from their pluralist egalitarianism. We sum up briefly these resulting preference relations.

A. (1) *Inequality under Certainty* versus (2) *Equal Risk, Unequal Final Well-being*.

(2) is morally preferred to (1) because it gives an equal shot of being cured for Ann and Bea whereas (1) cures one of the two arbitrarily.

B. (3) *Equality under Risk* versus (1) *Inequality under Certainty*.

(3) is morally preferred to (1) because not only does it give an equal shot of being cured for Ann and Bea but it also eliminates all interpersonal unfairness

but still maximizes total expected utility.

C. (3) *Equality under Risk* versus (2) *Equal Risk, Unequal Final Well-being*.

(3) is morally preferred to (2) for the same reasons in B.

D. (4) *Equality under Certainty* versus (4) *Inequality under Certainty*.

Suppose there is no cost to remove the inequality, say, for $c = 0$. In this case, (4) is morally preferred to (1) because inequality is suppressed without loss in expected total well-being.

E. (4) *Equality under Certainty* versus (2) *Equal Risk, Unequal Final Well-being*.

Suppose $c = 0$. (4) is morally preferred to (2) for the same reasons as in D.

F. (4) *Equality under Certainty* versus (3) *Equality under Risk*.

Suppose $c = 0$. (4) is morally indifferent to (3) because both offers Ann and Bea equal expected well-being and both are equally good prospects for each person and neither contains any inequality.

G. (4) *Equality under Certainty* versus (1) *Inequality under Certainty*.

For some, sufficiently small $c > 0$, (4) is morally preferred to (3) because it eliminates inequality.

H. (4) *Equality under Certainty* ($c > 0$ but c very small) versus (3) *Equality under Risk*.

For all $c > 0$ (3) is morally preferred to (4) because neither contains any inequality but (3) contains more valuable prospects for each person.

I. (5) *Equal Uncertainty, Unequal Final Well-being* versus (2) *Equal Risk, Unequal Final Well-being*.

(2) is morally preferred to (5) because this latter, due to the presence of uncertainty, reduces the value of individuals' prospects.

J. (6) *Equality under Uncertainty* versus (3) *Equality under Risk*.

(3) is morally preferred to (6) because (6) reduces the value of each individual's prospects and population level value prospect. (6) generates population-level uncertainty, because the decision-maker is uncertain about the anonymized distribution of final well-being⁵.

K. (7) *Unequal Uncertainty* versus (8) *Equal Moderate Uncertainty*.

⁵Hence, (5) does not generate population-level uncertainty but comes at the expense of certain inequality in the final value of well-being between Ann and Bea.

(8) is morally preferred to (7) because it distributes equally ex ante an equal quantity of uncertainty over Ann and Bea and thus its moral cost is shared equally among them.

L. (5) *Equal Uncertainty, Unequal Final Well-being* versus (1) *Equality under Certainty*.

For any decision-maker willing to incur a cost ($c > 0$) to eliminate the uncertainty and/or the inequality, (1) is morally preferred to (5).

Despite leading to the same outcome in terms of well-being (either cured or not), the risky and uncertain treatments do not bear the same moral values. The uncertain treatment is morally more costly than the risky treatment because it would expose Ann and Bea to experiencing unnecessary uncertainty, which is morally impermissible. Overall, uncertainty aversion and inequality aversion incur a cost to remove inequality. Hence, if the decision-maker is inequality averse but uncertainty neutral then for her, (2) and (3) are equivalent.

§2 Discussion

2.1 Some Questions

In this section, I point out some directions we could further discuss in the comments. Before drawing out one objection below, I rather prefer to suggest four questions to discuss:

1. Is it always reasonable to display uncertainty aversion for societal decisions even if the individual-level and population-level value prospects of well-being are very high?
2. For any distributive theory of justice under severe uncertainty (or ambiguity), should we consider in priority the experience of uncertainty as an impermissible moral cost and avoid it at all expenses, even if we could violate rational principles or not?
3. when a social planner has no information regarding citizens' attitudes towards uncertainty, which option should she take between a severely uncertain option with high pay-offs
4. when a social planner has information regarding citizens' attitudes towards uncertainty, does she have the legitimacy to override this information and always follow the recommendation made by an uncertainty-averse distributive theory of justice?

2.2 An Objection

In the following, I consider the theme of these questions in one precise case. Consider the following case. You are the decision-maker (DM) for Ann. Two treatments are at your disposal to cure Ann. Each treatment is decomposed into two parts and both steps have to be successful for the treatment to be effective overall and thus to cure Ann fully. Both treatments are sequential (the first part is followed in time by the second one). A treatment is effective for Ann if and only if both parts of the treatment are successful. As a representation of this decision problem, we adapt a recent example provided by Fleurbaey (2019⁶): the treatment is effective for Ann (at different levels of individual's value prospects) if you pick two balls of the same color. Let's call these two treatments, "option 9" and "option 10".

- (9) Option 9: you pick a ball from a risky urn, U_1 and observe its color. It is either Red with 50% of chance or Black with 50% chance. Then you decide between two options: either you to pick from an ambiguous urn U_2 (option 9a) or from the same risky urn⁷, U_1 (option 9b).
- If you choose option 9a, then there are possible outcomes.
 - You draw a ball with the same color from U_1 , then Ann is cured and she ends up with a lifetime well-being of 80.
 - Ann is not cured at all otherwise.
 - If you choose option 9b, then there are two possible outcomes
 - You draw a ball with the same color from U_1 , then Ann is cured and she ends up with a lifetime well-being of 50.
 - Ann is not cured at all otherwise.
- (10) Option 10: you pick a ball from the ambiguous urn, U_2 , and then necessarily from the risky urn, U_1 . There are two possible outcomes.
- If you draw two balls of the same color from U_1 and U_2 , then Ann is cured and she ends up with a lifetime well-being of 60.
 - Ann is not cured at all otherwise.

As a reminder, according to pluralistic, uncertainty-averse egalitarianism, you should be uncertainty averse. However, from this example above, it is not clear which option you should choose. There seems to be two ways to be uncertainty averse⁸. One way seems clearly irrational and another way, seems as Fleurbaey calls, "unreasonable".

⁶Manuscript. See the original examples provided in Al-Najjar and Weinstein 2009

⁷you put the previous you drew back in the urn before drawing this ball

⁸If you are neutral to uncertainty aversion, then you focus on the outcome's value and thus you choose option 9 over option 10 and then you choose option 9a over option 9b.

First let see the irrational way to uncertainty averse, which pluralistic uncertainty-averse egalitarianism could advice you but clearly you should not follow. This way corresponds to being averse to uncertainty and a non-expected utility maximizer. In this context, you can be either a naive non-expected utility maximizer or a sophisticated non-expected utility maximizer.

On the one hand, if you are a naive decision-maker then you compare the first step of option 9 and the option 10. Since, following pluralistic egalitarianism, the experience of of severe uncertainty should count as a moral cost, you discard option 10 and prefer the first step of option 9. since it is risky. You choose option 9 and you draw a ball from the risky urn, U_1 and you end up either with a red or a black ball. Then, you learn that you have two options that you had not foreseen, option 9a and option 9b. Option 9a is just as maximally uncertain as the first step of option 10 that you already rejected. So you switch towards option 9b, which is dominated by 10. If you had thought twice, you would had been better off by going directly at first with option 10, but you couldn't because you followed the uncertainty aversion principle and had limited rationality. This irrationality can be viewed as dynamic inconsistency.

On the other hand, if you are sophisticated decision-maker, you anticipate the second step of option 9 (the choice between 9a and 9b) and you depress the value of option 9a because of the presence of ambiguity. And option 9b is dominated by option 10. Thus you go directly with option 10 where finally you have 50% chances to have two colors of the same urn (since the ambiguous urn is followed automatically by a risky urn, the ambiguity in the first round has no negative effect on you). You cure Ann but not fully. Despite not being dynamic inconsistent, the sophisticated agent violates the Independence of Irrelevant Alternatives.

For these reasons, the moral permissibility of a distributive theory of justice designed for society as a whole (and not for each individual), based on the violations of such rationality principles, should be rejected. This could be the original recommendation emerging from Rowe and Voorhoeve's use of the Hurwicz criterion, independently from expected utility theory.

Second, let see the unreasonable way to be uncertainty averse, which pluralistic uncertainty-averse egalitarianism could also advice you. This way is compatible with being an expected-utility maximizer. In this context, you will choose option 10 over option 9b because the latter is dominated by option 10. Besides, the cost of experiencing uncertainty in 9a is so high that its expected value would be lower than the expected value of option 10, therefore you go with option 10. This is the option that pluralistic egalitarianism that recommends, by embedding the Hurwicz criterion with expected utility⁹. This

⁹One might look at Gul and Pessendorfer (2015) and Binmore (2016) for such kind of a model: the main difference with the original Hurwicz's criterion is that here the concept of uncertainty is decomposed as *perceived uncertainty* and *source uncertainty attitude*. The latter corresponds to the original Hurwicz's criterion that Rowe and Voorhoeve integrate in their definition of equality under uncertainty but the

alternative might be unreasonable compared to option 9a.

If the decision maker has access to Ann's information with regards to her attitude and she is willing to go with 9a, then there is no reasonable justification, as for me, to overcome her attitude by imposing a cautious attitude from the social planner on Ann and thus depriving her from an extra lifetime well-being (gained from 9a). From an aggregation perspective, the correct uncertainty attitude that the social planner should take into account is the one emerging from the citizens, except if she has another better - qualitatively speaking - source of information. But in any event, imposing one's own attitude for the entire society seems, at least, questionable and the same for all public policy decisions, even more (evidence might show that citizens are uncertainty averse towards specific public decisions and uncertainty seeking towards others).

From this short discussion, one take away is that the Hurwicz criterion, embedded in expected utility approach (let's note it "HEU") seems to be flexible enough to accommodate different levels of uncertainty attitudes depending on the context of risk and severe uncertainty. For instance, if maximally uncertainty averse HEU agents ($\alpha = 1$) always prefer bets on less uncertain sources, HEU agents with an intermediate level of uncertainty aversion ($0 < \alpha < 1$) reverse this preference when the probability of winning is low (see Gul and Pesendorfer 2015).

former seems not mobilized and integrated in the value functions of their pluralistic egalitarianism.