Ambiguity aversion behind the veil of ignorance

H. Orri Stefánsson^{1,2}

Received: 5 February 2019 / Accepted: 26 October 2019 / Published online: 15 November 2019
© The Author(s) 2019

Abstract

The veil of ignorance argument was used by John C. Harsanyi to defend Utilitarianism and by John Rawls to defend the absolute priority of the worst off. In a recent paper, Lara Buchak revives the veil of ignorance argument, and uses it to defend an intermediate position between Harsanyi's and Rawls' that she calls Relative Prioritarianism. None of these authors explore the implications of allowing that agent's behind the veil are sensitive to ambiguity. Allowing for *aversion* to ambiguity—which is both the most commonly observed and a seemingly reasonable attitude to ambiguity—however supports a version of Egalitarianism, whose logical form is quite different from the theories defended by the aforementioned authors. Moreover, it turns out that the veil of ignorance argument neither supports standard Utilitarianism nor Prioritarianism unless we assume that rational people are insensitive to ambiguity.

Keywords Veil of ignorance · Distribution-sensitive utility · Risk · Ambiguity

1 Introduction

The *Veil of Ignorance* is a powerful tool, that has been used by different authors to defend divergent views in distributive ethics. The function of the veil is to blind people to both their social positions as well as to their endowments and attitudes, and is meant to ensure that the preferences that people reveal between different “social gambles” are indicative of—or, alternatively, determine—the just arrangement of social institutions and the fair distribution of goods. The thought is that if people neither know what roles they occupy nor what assets or attitudes they have, then their preferences will be equally sensitive to the interests of all, and will therefore correspond to those of a fair-minded social-planner.

The term, “veil of ignorance”, was coined by Rawls (1971), who famously used it to defend the view that benefits to the worst off take absolute priority over benefits to

✉ H. Orri Stefánsson
orri.stefansson@iffs.se
<http://www.orrstefansson.is>

¹ Department of Philosophy, Stockholm University, Stockholm, Sweden

² Institute for Futures Studies, Stockholm, Sweden

the better off. But before Rawls coined the term, Harsanyi (1953, 1955) had used a similar construction to defend (Average) *Utilitarianism*, according to which the value of benefiting a person depends only on the size (in utility) of the benefit, and average welfare should be maximized.¹

In a recent article, Buchak (2017) revives the idea of a veil of ignorance to defend a theory she calls *Relative Prioritarianism* (often called Rank-Weighted Utilitarianism), and which is an intermediate between the theories defended by Harsanyi and Rawls. Drawing on her recent theory of rational attitudes to risk (Buchak 2013), she sides with Rawls in claiming that rational preferences behind the veil of ignorance support the view that benefiting a person matters more the worse off she is. But she also agrees with Harsanyi's view that all benefits—including those to the very best off—always count for something.

Rawls' argument differs from Buchak's and Harsanyi's arguments in whether people behind the veil are taken to know the probability with which they occupy the different social positions: whereas Rawls stipulated that people behind the veil have no such probabilistic information, Buchak and Harsanyi stipulated that the people know precisely these probabilities. The approach explored in this paper could be seen as a compromise between these two extreme positions: It will be stipulated that people behind the veil may not know the precise probability with which they occupy the different social position, but it will be assumed that they have some beliefs about these probabilities.²

The assumption that people behind the veil have some, but not precise, probabilistic information, makes it possible to ask what attitudes they take to what is often called *ambiguity*, in particular, whether they prefer known probabilities (or "chances") over unknown ones. Inspired by Buchak's (2017) argument that we should take people behind the veil to adopt the most risk averse attitude within reason, I shall argue, first, that ambiguity aversion is a form of risk aversion,³ and, consequently, that we should assume that people behind the veil are averse to ambiguity.

¹ All existing attempts to apply decision theory in the veil of ignorance argument replace a *fixed* number of states of the world (or events) with a *fixed* number of individuals (or groups of individuals). Hence, although these arguments strictly speaking support some average view, be it Utilitarian, Prioritarian, or Egalitarian, average and total views are extensionally equivalent given the employed framework. Alternative applications are possible, and would certainly be worth carrying out. In this paper I shall nevertheless follow the tradition in assuming a fixed state space and, correspondingly, a fixed population. And since average and total views are extensionally equivalent given such a framework, the prefix "average" will typically be omitted.

² Mongin's (2001) generalization of the veil of ignorance argument can similarly be seen as a compromise between the epistemic assumptions made on one hand by Buchak-Harsanyi and on the other hand by Rawls. Like me (and unlike Rawls), Mongin assumes that people behind the veil have some (subjective) estimates of the likelihood of corresponding to each member in the ranked distributions. Moreover, Mongin assumes, like me (and unlike Buchak and Harsanyi), that no probability assignment is forced upon people behind the veil. However, Mongin assumes that people's preferences between gambles with objective probabilities satisfy the von Neumann-Morgenstern (1944) axioms, and that people's preferences between gambles with subjective probabilities satisfy the additional Anscombe-Aumann (1963) axioms. In contrast, I shall only assume that people's preferences satisfy the Jeffrey-Bolker (1965, 1966) axioms. (Another author who (very briefly) explores ambiguity aversion behind the veil is Ellsworth 1978.)

³ The claim that ambiguity aversion is a form of risk aversion is defended by Bradley (2016) and Stefánsson and Bradley (2019).

The main aim of this paper is to examine what theory of distributive ethics is supported if one assumes that people behind the veil of ignorance are both ambiguity averse and have the so-called *Allais preference* (Allais 1953), which Buchak's theory is designed to accommodate. To this end, I apply to the veil of ignorance setting a theory that was recently developed by Stefánsson and Bradley (2015, 2019) specifically to simultaneously account for ambiguity aversion and Allais-type risk aversion. The resulting argument then supports a form of Egalitarianism that I call *Distribution-Sensitive Utilitarianism*,⁴ and whose logical form is quite different from the theories defended by Buchak, Harsanyi, and Rawls. Moreover, we'll see that the veil of ignorance argument supports neither *standard* (i.e., additively separable; see fn. 25) Utilitarianism nor Prioritarianism *unless* we assume that people behind the veil are insensitive to ambiguity.

If one accepts the normative importance of the veil of ignorance argument, then there are (at least) three different lessons for distributive ethics that one could draw from this paper. First, those who are committed to the rational permissibility and reasonableness of sensitivity to ambiguity can take the present paper to be an argument in favor of Distribution-Sensitive Utilitarianism. Second, those who don't like (standard) Utilitarianism and Prioritarianism can read the paper as an argument that rational people behind the veil can be sensitive to ambiguity. Third, those who like (standard) Utilitarianism and Prioritarianism, can read the paper as an argument that rational people behind the veil are *not* sensitive to ambiguity.

In contrast to the interpretations suggested above, one could also read this paper as undermining the normative importance of the veil of ignorance argument. For instance, Prioritarians and Utilitarians who are convinced that ambiguity sensitivity is rationally permissible, can read the paper as a *reductio* of one of the main premises of any veil of ignorance argument, namely, the premise that the preferences of people behind the veil reveal or determine principles of distributive ethics. Alternatively, such scholars could question the veil of ignorance framework employed in this paper, in particular, they could argue that we should assume knowledge of precise probabilities behind the veil.⁵ Finally, one could read this paper as showing, in conjunction with the arguments of Harsanyi, Rawls and Buchak, that since the veil of ignorance argument is so sensitive to subtle modelling choices, the argument does not settle the debates between the main competing views in distributive ethics.

⁴ People might object to the term, Distribution-Sensitive Utilitarianism, on two grounds. First, they might object to it on the grounds that *all* Utilitarianism is distribution sensitive, for instance, since consumption and wealth has decreasing marginal utility. As we shall see, however, the distribution to which the theory in question is sensitive, is a distribution of individual welfare (so not just consumption or wealth). Second, and in light of the response to the first objection, some might object that Distribution-Sensitive Utilitarianism is an oxymoron: by definition, they might claim, Utilitarianism is not sensitive to distribution of welfare. In response, it should be emphasized that "Utilitarianism", as the term is used here, simply refers to the theory that utility, understood as a numerical (typically, but not necessarily, cardinal) representation of a relation should be maximized. But that leaves open how utility is determined, and is, as we shall see, consistent with various forms of Egalitarianism being classified as Utilitarianism.

⁵ I thank a referee for suggesting this reading.

2 The veil of ignorance argument

In this section, I sketch the veil of ignorance argument as it will be employed in this paper. As most parts of the argument are very well-known, in particular, from the works of Harsanyi and Rawls, I shall focus on how the present argument differs from Harsanyi's and Rawls', as well as from Buchak's.

It will be assumed that people behind the veil of ignorance evaluate and rank *tuples of social groups* (i.e., lists of social groups), where each group can be thought of as an equivalence class of welfare.⁶ Intuitively, it makes sense to think of these tuples as welfare distributions, and they will frequently be referred to as welfare distributions throughout the paper. Note, however, that welfare, thus understood, is not (yet) a quantitative notion, but a comparative one, unlike the utilities that will later be introduced to represent the ranking of these welfare distributions. That is, these welfare distributions are not distributions of quantities, but rather distributions of groups of people where people within a group are deemed equally well-off. Moreover, the ranking can intuitively be thought of as a *preference* ranking, i.e., the idea is that people behind the veil ask themselves which distribution they would prefer to be actual, without knowing to which group in the distribution they will then belong.

The argument can be presented in five premises.⁷ First, following the tradition of Harsanyi and Rawls, it will be assumed that the preferences that people display behind the veil of ignorance correspond to the relative goodness of the ranked distributions, and (assuming teleology) thus inform us what distribution of welfare we should aim to achieve.

Second, interpersonal comparability of utility and welfare will be assumed. More precisely, I assume (with Buchak and Harsanyi) that individuals behind the veil of ignorance all agree on a unique⁸ *cardinal* (utility) representation of the welfare of the members of the ranked distributions. This is not to say that the individuals in the ranked distributions have a shared utility function. Rather, the assumption is that when two individuals behind the veil consider two individuals in a ranked distribution, the individuals behind the veil agree on a cardinal representation of the welfare of the two individuals in the distribution. Due to this and the fourth assumption, discussed below, we can treat choice behind the veil as the choice of a *single* individual. That is, when representing the attitudes of people behind the veil, only one utility function is needed. So, no bargaining or collective deliberation needs to be modeled or assumed, nor need we assume that any social aggregation procedure takes place behind the veil.

Third, it will be assumed, with Buchak and Harsanyi, that people behind the veil can form some (meaningful) assessments of the probabilities of corresponding to different members of the distributions that they are evaluating. But unlike Buchak and Harsanyi, I allow for the possibility that people behind the veil are not fully confident in their probability judgments. More precisely, it will be assumed that people behind the veil assign subjective probabilities to different objective probability (or “chance”)

⁶ This means that we assume that the veil of ignorance argument concerns only how welfare should be distributed (and not which rights and liberties we should grant people), which makes the argument more limited than the version developed by Rawls.

⁷ The presentation closely follows Buchak's (2017).

⁸ That is, unique up to a point of scale and starting point.

distributions that specify the chance of corresponding to each member, but (contra Buchak and Harsanyi) it will be assumed that they may assign a positive subjective probability to more than one such distribution. As previously mentioned, this can be seen as a compromise between the positions of Harsanyi and Buchak, on the one hand, and on the other hand the position of Rawls, who assumed that no such probability assignments can be made behind the veil. The main motivation for this third assumption is to explore what implications this particular compromise has for distributive ethics.

Fourth, I shall follow Buchak in assuming that when choosing for others whose risk attitudes one does not know, one should “err on the side of caution and choose the less risky option, within reason” (Buchak 2017, p. 631). To support this claim, Buchak (330–336) considers what risks one could justifiably take if one stumbled upon a person in need of one of two medical interventions that carry with them different levels of risk. In such a situation, she suggests, it seems intuitive that, other things being roughly equal, one would be morally required to choose the less risky of the two interventions if one did not know the person’s risk attitudes.^{9,10} Applied to the veil of ignorance argument, this intuition supports ascribing to individuals behind the veil of ignorance “the most risk-avoidant reasonable risk attitude” (638).¹¹ (Note that the potentially diverging risk attitudes and beliefs of the individuals in the ranked distributions are however ignored.¹²)

Fifth, it will be assumed that aversion to ambiguity can be a perfectly reasonable form of risk aversion; thus, I shall assume that people behind the veil can be ambiguity averse (as well as having the Allais preference). Since the assumption that ambiguity aversion is (or can be) a form of risk aversion is the most novel part of the present version of the veil of ignorance argument, I devote special attention to it in Sect. 3.

⁹ If one however knows a person’s risk attitude, then she thinks one should defer to it.

¹⁰ It is worth noting, as a referee points out, that by “risk” attitudes Buchak does not mean only the curvature of the person’s utility function (as has become tradition in, for instance, economic theory). Rather, she means a more general form of risk aversion that is captured by, for instance, her risk-weighted expected utility theory.

¹¹ Although Buchak does not explicitly say so, it seems that she takes the reasonable attitudes behind the veil to be a subset of the set of rational attitudes. I shall not make that assumption, however, but will instead leave open the possibility that an attitude could be reasonable without being rational. After all, it would seem that an attitude could be reasonable without being rational in the strict decision-theoretic (or coherence) sense. I shall get back to the concepts of reason and rationality in later sections.

¹² Well-known problems arise when taking people’s heterogeneous beliefs into account in social evaluations. In particular, there is a conflict between heterogeneity in beliefs at the individual level and seemingly natural constraints on the social evaluation. (See e.g. Mongin 1995; Bradley 2005 for discussions of problems of this kind when it comes to preference aggregation, and Mongin 2001 for a discussion of analogous problems for the veil of ignorance argument.) These problems can be avoided by relaxing either the assumptions of individual rationality (to some extent in the line with the relaxation that I suggest below), or by relaxing the constraint on the social evaluation. (For an overview of this literature, as well as a sophisticated new approach to aggregation under uncertainty, see Mongin and Pivato ms.) Nevertheless, to keep things as simple as possible, I shall from now on assume that the beliefs of the individuals in the ranked distributions are ignored, in the sense that only what we could think of as people’s *actual* welfare counts. Thus, it will be assumed that people’s beliefs are taken into account in the social evaluation only in so far as they *directly* affect people’s welfare (e.g. by making them happy or upset). In contrast, people’s potentially misguided beliefs about how their interests are best served are ignored.

To capture the above assumptions about the attitudes of people behind the veil, I apply, as previously mentioned, a theory recently developed by Stefánsson and Bradley (2015, 2019), and find that the argument then supports the aforementioned Distribution-Sensitive Utilitarianism.

3 Aversion to ambiguity

In this section I explain sensitivity to ambiguity, that is, an agent's sensitivity to how uncertain she is about the relevant probability distribution. The importance of attitudes to ambiguity was originally brought to the attention of decision theorists and economists through the so-called *Ellsberg paradox* (Ellsberg 1961), one version of which is presented in the Appendix (A.1). But they can be illustrated much more simply with the following example (Stefánsson and Bradley 2019).¹³

Suppose that you have in front of you a coin, C_1 , that you know to be perfectly symmetric, and that you know to have been tossed a great number of times and has come up heads exactly as many times as it has come up tails. More generally, suppose that you possess the best possible evidence for the coin being unbiased. Two questions: (1) To what degree do you believe that C_1 will come up heads on the next toss? (2) How much would you be willing to pay for a gamble whereby you get \$10 if C_1 lands heads up on its next toss but get nothing otherwise?

Now suppose instead that you have in front of you a coin, C_2 , that you know to be either double headed or double tailed, but you don't know which way it is biased. Again, two questions: (1') To what degree do you believe that C_2 will come up heads on the next toss? (2') How much would you be willing to pay for a gamble whereby you get \$10 if C_2 lands heads up on its next toss but get nothing otherwise?

Many people seem to use something like the Principle of Insufficient Reason¹⁴ when answering questions like (1') (see e.g. Binmore et al. 2012). Since they have no more reason for thinking that the coin will come up heads than tails, they are equally confident in these two possibilities. So, since these possibilities (let's assume) exhaust the possibility space, they should believe to degree 0.5 that the second coin comes up heads. But that is, of course, the same degree to which they should believe that the first coin comes up heads, assuming something like Lewis' (1980) Principal Principle, which, informally, states that one's knowledge about chances, i.e., objective probabilities, should guide one's degree of belief, i.e., subjective probabilities.

¹³ It is worth emphasizing that although the simple illustration assumes that degrees of belief can be measured independently of preference, many experiments on ambiguity sensitivity do not make that assumption. In fact, Ellsberg's (1961) suggested experiments, one of which is discussed in Appendix A.1, make no such assumption.

Another difference between on one hand the simple illustration and on the other hand the Ellsberg paradox, is that the former but not the latter employs two distinct probability spaces. (Thanks to Philippe Mongin for bringing this to my attention.)

¹⁴ Roughly speaking, the Principle of Insufficient Reason states that for any two propositions, if you have no reason for being more confident of the truth of one proposition than the other, then you should be equally confident in their truth. The principle can give rise to well-known problems when applied to sample spaces for which there are different but equally "natural" (finest) partitions. But these problems need not worry us in the present context, where there does seem to be a single most natural partition, namely {double headed, double tailed}.

What about questions (2) and (2')? A number of experimental results on Ellsberg-type decision problems have shown that a large share of subjects as diverse as students, trade union leaders, actuaries, and executives, are ambiguity averse, meaning that other things being equal, they prefer gambles with known chances of outcomes over gambles with unknown chances (for an overview of these experimental results, see e.g. Machina and Siniscalchi 2014; Trautmann and van de Kuilen 2015). More generally, people tend to prefer less spread in subjectively possible chances over more spread, other things being equal; i.e., they prefer the chance values that they don't take to be ruled out by their evidence to be spread over a smaller interval rather than a larger one, other things being equal. In the example under discussion, ambiguity aversion translates into a preference for a gamble on C_1 —where the subjectively possible chances are confined to a single value, 0.5, and thus minimally spread—over a gamble on C_2 —where the subjectively possible chances are spread over the whole zero-one interval—and hence a willingness to pay more for the first gamble than the second.¹⁵

The above gambles can be straightforwardly turned into “social gambles” or distributions like those that people behind the veil of ignorance are asked to rank. Suppose you know that social gamble, or distribution, d_i , confers a 0.5 chance on you leading an excellent life and a 0.5 chance on you leading a life that is barely worth living. Gamble d_j , in contrast, either condemns you for sure to a life barely worth living or ensures that you will lead an excellent life; but you don't know towards which life the gamble is biased and you find each bias to be equally likely. Then, if you are ambiguity averse, you prefer d_i over d_j .

This is not the place to argue for the rational permissibility of ambiguity aversion, as exemplified by a preference for the gamble on coin C_1 over the gamble on coin C_2 .¹⁶ For the present purposes, the important point is the claim that ambiguity aversion is a form of risk aversion (as argued in Bradley 2016; Stefánsson and Bradley 2019). In decision theory and economics, a person is said to be risk averse with respect to some (real valued) good $\mathcal{G}$ if she prefers a gamble g_i to another gamble g_j which has the same expected benefit of $\mathcal{G}$ as g_i but a greater spread in the possible values of $\mathcal{G}$. In other words, a person is risk averse if she prefers a gamble to a mean-preserving spread of it.

Now let's compare the above understanding of what it is to be risk averse with ambiguity aversion, as exemplified by a preference for a gamble on coin C_1 over a gamble on coin C_2 . A person with this attitude prefers one gamble g_i to another g_j which has the same expected chance of resulting in a good (winning \$10) as g_i but has

¹⁵ A perhaps seemingly alternative explanation of the attitudes under discussion, which nevertheless turns out to be logically equivalent to the above one given the framework presented in the next section, is that even *ex post*, people prefer to actually have had a chance, rather than just having thought that they had a chance.

¹⁶ But see e.g. Gilboa et.al. (2009) for an argument that it is, in some circumstances, *more* rational to have ambiguity averse preferences than ambiguity neutral ones. Examples of authors who take ambiguity aversion to be rationally permissible include Binmore (2008), Rowe and Voorhoeve (2018) and Stefánsson and Bradley (2019). Not everyone agrees, however, that ambiguity aversion is rationally permissible. Al-Najjar and Weinstein (2009) for instance argue that ambiguity aversion is generally irrational, while Fleurbaey (2018) argues that ambiguity aversion may be rationally permissible when it comes to personal preferences but not when it comes to social evaluations.

a greater spread in the chance of resulting in the good:¹⁷ the whole zero-one interval compared to just 0.5. In other words, this person prefers a gamble to a mean-preserving spread of it. So, we should say that this person is *risk-averse with respect to chances* of winning \$10.

Why is it important to establish that ambiguity aversion is a form of risk aversion? The reason is that it shows that if we accept the fourth premise of the veil of ignorance argument, then we should assume that people behind the veil adopt ambiguity averse preferences. Recall that this premise, which is adopted from Buchak, states that when choosing for someone without knowing the person's risk attitude, one should assume the most risk averse attitude within reason. This Buchak interprets as a rather strong requirement: "if a reasonable person *could reject* [an option on the grounds that it is too risky], then we are not justified in choosing it" on someone else's behalf (Buchak 2017, p. 631; emphasis added). Moreover, she (correctly, it seems) takes this to imply that people behind the veil should (and thus will) adopt the most risk averse preference within reason. So, if ambiguity aversion is a form of risk aversion, as I contend, and since on the face of it reasonable people (college students, union leaders, actuaries, executives, etc.) often do reject options on the grounds that they are too ambiguous, then it follows from the fourth premise that people behind the veil should (and thus will) adopt the most ambiguity averse preference within reason. Hence, people behind the veil will, other things being equal, opt for less ambiguous distributions.

Allowing for sensitivity to ambiguity can, as previously noted, be seen as a way of reaching a compromise between Harsanyi's and Rawls' epistemic assumptions about the veil of ignorance. Whereas Harsanyi effectively assumed that people behind the veil of ignorance would be fully confident in one particular probability distribution over the different social positions, namely, the uniform distribution, Rawls famously argued that people behind the veil have no basis for making such judgments and thus cannot form meaningful probability estimates for the different social positions. In contrast, I shall assume that people behind the veil use their beliefs about chances to form subjective estimates of probabilities, but need not be fully confident in these estimates, since they may consider multiple different objective probability distributions to be possible given each choice. Moreover, the thought that people may be sensitive to ambiguity behind the veil will be captured by assuming that, even holding fixed the subjective expectation of chance, people need not be indifferent between, on one hand, social gambles where the range of chance distributions that they take to be epistemically possible is narrow, and, on the other hand, social gambles where the range of epistemically possible chance distributions is wide.

¹⁷ As a referee points out, this assumes that the decision-makers of interest, while not having precise subjective probabilities, has judgments about different potential (objective) probability distributions. There are even "more ambiguous" situations that do not fit this description, but they will be set aside in this paper, where it will be assumed that, for instance, people behind the veil have judgments about different possible objective probabilities.

4 Incorporating ambiguity sensitivity

The decision theory that Buchak applies to her veil of ignorance argument, that is, her “risk weighted expected utility theory”, cannot capture ambiguity aversion (for a discussion, see Stefánsson and Bradley’s 2019). Nor, of course, can the theory that Harsanyi applied, that is, expected utility theory as developed by von Neumann and Morgenstern (1944). However, unlike the theory that Harsanyi applied, Buchak’s theory can account for the so-called *Allais preference* (named after Allais 1953). I shall not discuss this preference in detail, but the property of the Allais preference that I focus on—a property that is inconsistent with von Neumann and Morgenstern’s theory—is that it values, say, a n gain in the probability of winning a great prize differently if the prior probability was $1 - n$ than if the prior probability was $m < (1 - n)$.

The aim of this section is to explore what happens if we assume that people behind the veil of ignorance both have the types of risk attitudes that Buchak’s theory is meant to capture—in particular, Allais preference—and in addition are averse to ambiguity, as described in the last section. As far as I am aware, the only *explicitly normative* decision theory satisfying this constraint is Stefánsson and Bradley’s (2015, 2019) extension of Jeffrey’s (1965) theory to chance propositions. In this section a sketch of the theory in question is presented, but in the next section we shall see why, when employed in the veil of ignorance argument, the theory supports Distribution-Sensitive Utilitarianism.

To allow for ambiguity aversion, understood as suggested in the last section, we need a framework where both objective probabilities (or *chances*) and subjective ones (or *degrees of belief*) are represented.¹⁸ To achieve this, Stefánsson and Bradley allow agents to take both conative and cognitive attitudes to both factual propositions and chance propositions,¹⁹ defined as follows. Let Ω be a Boolean algebra²⁰ of propositions (i.e., sets of possible worlds) describing “ultimate” outcomes, e.g. that one wins \$10 (in an ordinary choice scenario) or lives at a particular level of welfare (in the veil of ignorance choice scenario). Let Δ be a Boolean algebra of propositions describing all possible objective probability (or chance) distributions over the propositions in Ω . For instance, for any proposition $X \in \Omega$ and any value $n \in [0, 1]$, there is a proposition $Ch(X) = n \in \Delta$ which is true just in case the chance of X is n (denoted $ch(X) = n$). Strictly speaking, we should think of these chance propositions as being time-indexed (since chances evolve), but as the veil of ignorance decision-problem is static, the time-index can be safely ignored.

Let $\langle Ch(X_i) = \alpha_i \rangle$ denote the intersection of the propositions $Ch(X_1) = \alpha_1$, $Ch(X_2) = \alpha_2$, ..., $Ch(X_n) = \alpha_n$, where $\{X_i\}_{i=1}^n$ is an n -fold partition of Ω and the α_i are real numbers in $[0, 1]$ such that $\sum_{i=1}^n \alpha_i = 1$. So, $\langle Ch(X_i) = \alpha_i \rangle$ is a gamble that

¹⁸ For the present purposes, objective probability is whatever plays the role of chance in the Principal Principle (discussed below). So, objective probabilities are any probabilities such that if you know them, then they should constrain your subjective probabilities. Hence, relative frequencies could count as objective probabilities (even in a deterministic world).

¹⁹ The possibility of agents taking (non-derivative) conative attitudes to chances is a crucial difference between on one hand Stefánsson and Bradley’s framework and on the other hand the framework of Anscombe and Aumann (1963), in which both subjective and objective probabilities are also represented.

²⁰ That is, a set closed under the Boolean operators.

results in outcome X_i with chance α_i . In other words, gambles like $\langle Ch(X_i) = \alpha_i \rangle$ correspond to lotteries in the von Neumann and Morgenstern (1944) framework.

Finally, let $\Gamma = \Delta \times \Omega$, and suppose that a rational person's preferences between propositions in Γ , and similarly between propositions in both Ω and Δ , satisfy the Jeffrey–Bolker axioms (Jeffrey 1965: ch. 9; Bolker 1966).

In the veil of ignorance setting, the above assumptions mean that the people are faced with a set Ω of possible (ordinal) welfare levels, formulated as propositions that form a Boolean algebra (which means that the set includes disjunctions of chance propositions, which contain uncertainty about chance); finally they are; they are faced with a set of chance distributions over these welfare levels (representing the chances that they will have each welfare level), also formulated as propositions that form a Boolean algebra; and faced with the crossproduct of these two sets (which also forms a Boolean algebra). Moreover, the people's preferences order the elements of these three sets in a way that satisfies the Jeffrey–Bolker axioms.

The above assumptions entail that there is (by Bolker's 1966 theorem) a subjective probability measure P on Γ , and a utility function u on the same set except that the contradictory proposition has been removed, relative to which the person's preferences, over say Δ , can be represented as maximizing:

Desirability For any $\mathcal{G} = \langle Ch(X_i) = \alpha_i \rangle \in \Delta$ and any n -fold partition $\{Y_i\}_{i=1}^n$ of $\mathcal{G}$:

$$V(\mathcal{G}) = \sum_{i=1}^n u(Y_i) P(Y_i \mid \langle Ch(X_i) = \alpha_i \rangle) \quad (1)$$

To take an example, let $\mathcal{G}$ be the proposition that you win \$10 if a coin comes up heads but nothing otherwise, i.e., intuitively the proposition that you gamble on heads. We can partition this proposition into two possibilities according to whether the coin comes up heads or tails (for simplicity, we can assume that it cannot land and stay on its edge). Hence, Desirability entails that the desirability of $\mathcal{G}$ equals the utility of the gamble when the coin comes up heads, weighted by your conditional subjective probability for the coin coming up heads given $\mathcal{G}$, plus the utility of the gamble when the coin comes up tails, weighted by your conditional subjective probability for the coin coming up tails given $\mathcal{G}$.²¹

Now let's briefly see how the theory sketched above can simultaneously account for the Allais preference and ambiguity aversion (but see Stefánsson and Bradley 2019 for a more thorough discussion). The simple reason why it can account for both these attitudes, is that since chance distributions are included in the domain of the utility function, it can accommodate the idea that the added benefit of an increase in the chance of a good might depend on the prior chance; for instance, the idea that a 0.1 gain in the chance of a good is more valuable when the prior chance was 0 than when the prior chance was 0.4. As an illustration, suppose an s-shaped utility function over chances, concave over almost the whole zero-one interval but with an inflection point

²¹ As those who are familiar with Jeffrey's framework will know, the utility function, u , is actually itself a desirability function. So, the "utility" of taking the gamble when the coin comes up heads also satisfies equation (1). Hence, the difference between the two functions u and V is nothing but their domain.

Fig. 1 S-shaped utility function (u) over chances (ch)

Fig. 2 Linear utility function (u) over money (\$)

at the very top after which the function is convex (as in Fig. 1), and a linear utility function over money (as in Fig. 2).

To see how a person with ambiguity averse preferences can be represented as maximizing desirability given the above types of functions, recall again the choice between a bet on coin C_1 and a bet on coin C_2 , where the first coin is known to be fair but the second is known to be either double headed or double tailed. Now if a person's utility function over chances is concave over most parts of the zero-one interval, then she will prefer a bet on the first coin over a bet on the second coin. Simplifying the notation somewhat,²² for such a person, it will be the case that:

$$u(Ch(\$10) = 1) - u(Ch(\$10) = 0.5) < u(Ch(\$10) = 0.5) - u(Ch(\$10) = 0) \quad (2)$$

In other words, for such a person, the undesirability of the risk that the second coin is double tailed outweighs the desirability of the possibility that the second coin is double headed, when compared to the first coin, a bet on which has a sure 0.5 chance of resulting in the prize. More generally, ambiguity averse preferences, w.r.t. a gamble G , can be represented by a utility function over G 's chances that is at least as concave

²² Since we know that we are dealing with gambles that either result in a “desirable” prize (\$10) or nothing (\$0), I only specify the chance of one of these outcomes for each gamble.

as the utility function over G 's possible prizes (Bradley 2016; Stefánsson and Bradley 2019).

The inflection point of the utility function over chances around the upper end of the zero-one interval, after which the function is convex, means that the utility function can also represent a person with the Allais preference. For instance, it can be seen from the graph that greater utility is gained by a increasing the chance of the prize in question by 0.1 when the prior chance was 0.9 than when the prior chance was 0.8.

5 Distribution-sensitive utilitarianism

I contend that rational people's preferences can in general be represented as maximizing the expectation of a Jeffrey value function, extended to chance propositions as suggested in the last section. What does this mean for the veil of ignorance argument? The answer is that the conclusion of the argument is a distributive theory that seems most natural to interpret as a form of Egalitarianism, but which I shall call *Distribution-Sensitive Utilitarianism*.

Before formulating this version of Egalitarianism, we need to reinterpret the gambles to make them suitable as objects of choice behind the veil of ignorance. Informally, we can now think of the X_i as different social groups (which might be as small as a single life), assuming that the different lives in each group are indistinguishable from the point of view of people behind the veil. More formally, we can interpret each X_i as an equivalence class of welfare (of lives), where these classes are determined by the relation $\lesssim$ on Ω , which is interpreted as the preference relation of people behind the veil. So, a numerical measure of welfare is not assumed from the start, but is entailed by the assumption that $\lesssim$ on Ω satisfies the Bolker-Jeffrey axioms.²³

Finally, we can reinterpret $Ch(X_i) = \alpha_i$ as a proposition specifying the chance of belonging to social group X_i . So, now think of $\mathcal{D} = \langle Ch(X_i) = \alpha_i \rangle$ as the distribution according to which the probability of belonging to social group X_i is α_i . And $u(X_i)$ is the utility of belonging to X_i , as determined by $\lesssim$ on Ω .

If the theory sketched in the last section correctly characterizes rational preference, then people behind the veil will choose the distribution $\mathcal{D}$ that, for any n -fold partition $\{Y_i\}_{i=1}^n$ of $\mathcal{D}$, maximizes:

$$DSU(\mathcal{D}) = \sum_{i=1}^n u(Y_i) P(Y_i \mid \langle Ch(X_i) = \alpha_i \rangle) \quad (3)$$

To see why it makes sense to think of the above quantity as distribution-sensitive (average; see fn. 1) utility, note first that $\{X_i \cap \langle Ch(X_i) = \alpha_i \rangle\}_{i=1}^n$ is an n -fold partition of $\langle Ch(X_i) = \alpha_i \rangle$. Second, I make the standard assumption of assuming that

²³ Note that $\lesssim$ denotes the (shared) preference relation of people behind the veil. The beliefs and preferences of the people in the ranked distributions are however not explicitly represented. I assume that $\lesssim$ ranks the lives in terms of how well-off the people leading them are judged to be, but I leave open the question of how welfare is determined and compared. (See fn. 12.)

rational people satisfy the Principal Principle, which in the present framework can be formalized as:²⁴

Principal Principle For any $\langle Ch(X_i) = \alpha_i \rangle \in \Delta$:

$$P(X_i \mid \langle Ch(X_i) = \alpha_i \rangle) = ch(X_i) = \alpha_i \quad (4)$$

In other words, the subjective probability one should assign X_i , if one is certain that the objective probability of X_i is α_i , is α_i .

Putting together the last two observations, (3) becomes:

Distribution-Sensitive (Expected) Utility For any $\mathcal{D} = \langle Ch(X_i) = \alpha_i \rangle \in \Delta$:

$$DSU(\mathcal{D}) = \sum_{i=1}^n u(X_i \cap \langle Ch(X_i) = \alpha_i \rangle) \alpha_i \quad (5)$$

In other words, rational preferences behind the veil support the distribution that maximizes distribution-sensitive utility, which, for any distribution $\mathcal{D}$, is calculated by, first, figuring out the utility u of belonging to each group when the distribution is $\mathcal{D}$, second, weighing u by the probability of belonging to that group, and, finally, adding up these probability weighted and distribution-sensitive utilities. Moreover, assuming general risk aversion behind the veil, Distribution-Sensitive Utilitarianism is distribution-sensitive in an Egalitarian way.

To take an example, suppose that the distribution we are evaluating contains two equiprobable social groups, X_1 and X_2 , and let the people in X_1 be very well off while the people in X_2 are not so well off. To calculate the distribution-sensitive utility of this distribution, (X_1, X_2) , we find the utility of belonging to group X_1 when the distribution is (X_1, X_2) , weigh that utility by 0.5, and add the result to the utility of belonging to group X_2 when the distribution is (X_1, X_2) weighted by 0.5. Since the measure is distribution-sensitive in an egalitarian direction, at least one of the utilities in this sum is lower than the corresponding utility when the distribution is (X_1, X_1) , i.e., when everyone is very well off.

To make this more concrete, suppose that $u(X_1) = 9$ and $u(X_2) = 4$, which we can interpret as, say, the utility of belonging to group X_1 is 9 when no information is added about the distribution of welfare levels. Intuitively, it makes sense to think of these numbers as isolating the individual welfare of belonging to each group when one ignores possible concerns one may have about the welfare enjoyed by others.

To simplify the notation, let's now use Ch to denote the above gamble, that is, an equal chance between belonging to X_1 and X_2 . Moreover, suppose that $u(X_1 \cap Ch) = 3$, $u(X_2 \cap Ch) = 2$. As we shall see below, the assumptions in this paragraph are consistent with the assumptions in the last paragraph just in case one does not assume *Chance Neutrality*, which standard Utilitarianism entails. While it made sense to interpret the first pair of numbers, 9 and 4, as the individual welfare of belonging to each group, it would make more sense to interpret the second pair of numbers, 3 and

²⁴ To simplify, I omit the much discussed and debated “admissibility” proviso (Lewis 1980).

2, as the *moral value* of some person belonging to each group when the distribution is Ch .

Now let's consider the distribution where everyone belongs to group X_1 , and, again to simplify the notation, call this distribution Ch^+ . The assumption of ambiguity aversion behind the veil entails that, say, $u(X_1 \cap Ch) < u(X_1 \cap Ch^+)$, which means that the value of a person belonging to X_1 is greater when everyone is at the same level of welfare than when half the population is at a lower level. In other words, by the veil of ignorance argument, the assumption of ambiguity aversion at the individual level entails inequality aversion at the social level. In contrast, Utilitarianism on the social level holds that $u(X_1 \cap Ch) = u(X_1 \cap Ch^+)$, which, as we shall soon see, is equivalent to ambiguity neutrality behind the veil.

As the above example illustrates, it would seem natural to interpret the instance of Distribution-Sensitive Utilitarianism that the present version of the veil of ignorance argument supports as Egalitarianism. In particular, if we follow the tradition (going back to Parfit 1991) of calling Egalitarian any theory that allows for the possibility that the value that a person contributes towards the overall value of a distribution is not fully determined by the person's welfare but is also partly determined (in an Egalitarian way) by the distribution of welfare, then Distribution-Sensitive Utilitarianism is Egalitarian.

It might be worth noting that Distribution-Sensitive Utilitarianism can account for moral benefits of equality that standard Utilitarianism cannot. There are two ways in which standard Utilitarianism can account for the value of equality. On the one hand, if we assume that consumption and wealth has decreasing marginal benefits, then an equal distribution of consumption will, other things being equal, result in greater aggregate welfare than an unequal distribution. On the other hand, if people care about equality, in the sense that inequality decreases their welfare, then an equal distribution of, say, consumption and wealth, will, other things being equal, lead to greater aggregate welfare than an unequal distribution. Hence, standard Utilitarianism will, other things being equal, favor an equal distribution of consumption and wealth over an unequal one, provided that consumption and wealth have decreasing marginal benefit, and/or individuals in the evaluated population care about equality.

Now, Distribution-Sensitive Utilitarianism can account for the above two *instrumental* ways in which equality might matter. But in addition, it allows for the possibility that an equal distribution of welfare is *intrinsically* valuable, in that its value is not exhausted by the value that equality contributes in terms of increasing individuals' welfare. To take an extreme example: Distribution-Sensitive Utilitarianism can account for the possibility that equality matters even when evaluating a population where *both* nobody is an egalitarian *and* consumption has constant marginal utility for everyone.

Distribution-Sensitive Utilitarianism is however consistent with various different (and potentially competing) ways in which the distribution might affect the contributive value of a life (or a group) at a particular welfare level. For instance, Distribution-Sensitive Utilitarianism is consistent with the idea that the contributive value of a life at a particular welfare level is partly determined by how many people have that level of welfare as compared to how many people are better and worse off. Moreover, Distribution-Sensitive Utilitarianism is consistent with the thought that the distribution effect on the contributive value of a life at a particular welfare level is determined by the distance in welfare between different groups. Finally, Distribution-Sensitive

Utilitarianism is of course consistent with any combination of the aforementioned two views. More generally, Distribution-Sensitive Utilitarianism is consistent with various competing Egalitarian views about how and why equality matters; all it says is that how welfare is distributed can matter to how much any particular welfare level contributes to overall value.

However, if we assume that rational people behind the veil of ignorance satisfy what Stefánsson and Bradley (2015) call *Chance Neutrality* (to be explained below), then the conclusion of the veil of ignorance argument is *standard*²⁵ Utilitarianism. So, what I have called Distribution-Sensitive Utilitarianism is not inconsistent with (standard) Utilitarianism; rather, the latter is a special case of the former. (Nor is Distribution-Sensitive Utilitarianism inconsistent with Prioritarianism, as we shall soon see.)

Informally, Chance Neutrality is the view that once it is known which outcome obtained, it is a matter of ethical and practical irrelevance what the chances were. As Stefánsson and Bradley (2015) note, Chance Neutrality entails what we could call *instrumentalism* about chances: chances only matter in so far as they make good or bad outcomes more or less likely, but they are of no value in and of themselves. To take an example, Chance Neutrality entails that if a patient dies of kidney failure, then the fact that he had, say, the same chance as any other patient of receiving the hospital's only kidney makes no difference. More formally:

Chance Neutrality For any $\langle Ch(X_i) = \alpha_i \rangle$ and any X_i :

$$u(X_i \cap \langle Ch(X_i) = \alpha_i \rangle) = u(X_i) \quad (6)$$

Now, adding Chance Neutrality to (5) gives us:

Average Utility For any $\mathcal{D} = \langle Ch(X_i) = \alpha_i \rangle$:

$$DSU(\mathcal{D}) = \sum u(X_i)\alpha_i \quad (7)$$

Recall that, for the social gambles, $u(X_i)$ is interpreted as the utility of belonging to group X_i while α_i is the probability of belonging to group X_i . So, if people behind the veil of ignorance satisfy Chance Neutrality, in addition to Desirability and the Principal Principle, then the distributive principle that the argument from the veil of ignorance supports is standard (average; recall fn. 1) Utilitarianism.

Stefánsson and Bradley (2015) argue at length that Chance Neutrality is not a requirement of rationality. In fact, it is the denial of Chance Neutrality that allows for a unified treatment of the Allais preference and ambiguity aversion (Stefánsson and Bradley 2019). Hence, if we want to allow for both types of attitudes behind the veil,

²⁵ “Standard” Utilitarianism is the view that the moral value of a distribution is a linear additively separable function of the (final) welfare of the individuals in the distribution. So, Distribution-Sensitive Utilitarianism is evidently not standard Utilitarianism. From now on, just “Utilitarianism” means standard Utilitarianism (as opposed to e.g. Distribution-Sensitive or Rank-Weighted Utilitarianism). Similarly, “standard” Prioritarianism is the view that the moral value of a distribution is a concave additively separable function of the welfare of the individuals in the distribution. From now on, just “Prioritarianism” means standard Prioritarianism. (Note that standard Utilitarianism and Prioritarianism are both what is called *ex post* views. That is, they are ultimately concerned with people's final, or actual, welfare, not their expected welfare.)

then we should not assume that people behind the veil of ignorance satisfy Chance Neutrality, and thus we should not assume that the distributive principle that their preferences support is Utilitarian. Instead, we should assume that it is distribution-sensitive.

It is important to note that we do not have to assume any particular risk attitude—for instance, we do not have to assume any particular degree of ambiguity sensitivity and risk aversion—to get the above result. That is, we do not have to assume some particular attitude to risk to derive Distribution-Sensitive Utilitarianism from the present version of the veil of ignorance argument. All we need to assume, to get that result, is that people behind the veil do not satisfy Chance Neutrality.

Since ambiguity plays a central role in the present paper, it is worth examining in more detail the relationship between Chance Neutrality and ambiguity neutrality (ambiguity neutrality being the negation of ambiguity sensitivity). For while Chance Neutrality and ambiguity sensitivity are both properties of attitudes to chances, it might not be obvious that they conflict. After all, Chance Neutrality concerns attitudes to chance distributions once it is known what outcome obtains, whereas ambiguity sensitivity is an attitude to chance distributions before it is known which outcome obtains.²⁶ Therefore, it is worth going through a concrete example that illustrates how Chance Neutrality rules out ambiguity sensitivity (given Desirability and the Principal Principle). But first, let's consider some general results that show the tension between Utilitarianism and ambiguity sensitivity.

To prove general results about the connection between ambiguity neutrality and Utilitarianism, we need a precise and general definition of ambiguity neutrality. Such a definition however turns out to be quite complex, so I leave it to the Appendix (A.1). But informally, the definition simply says that an ambiguity neutral person doesn't care if chances are spread or concentrated as long as her subjective expectation of chance stays the same. On the basis of such a definition, I prove in Appendix A.2:

Proposition 1 *If a person satisfies Desirability, the Principal Principle, and Chance Neutrality, then she is ambiguity neutral.*²⁷

From the point of view of distributive ethics, the important thing to note about the above result is that it entails that if we allow for sensitivity to ambiguity behind the veil, then the resulting distributive principle will not be Utilitarian.²⁸ For given Desirability and the Principal Principle (which, I contend, are both normatively unassailable), and assuming the veil of ignorance argument, it is straightforward to verify that Utilitarianism on the social level is *equivalent* to assuming Chance Neutrality on the individual level. And Proposition 1 tells us that we cannot assume Chance Neutrality if we want to allow for sensitivity to ambiguity.

The same is true of (standard; see fn. 25) Prioritarianism: If we allow for sensitivity to ambiguity behind the veil, then the resulting distributive principle will not be Prior-

²⁶ Thanks to Seth Lazar for encouraging this clarification.

²⁷ As observed in Appendix A.2, the importance of the Principal Principle for this result is rather limited, since only a slight modification (but, in my view, worsening) of the definition of ambiguity neutrality is required for Desirability and Chance Neutrality to entail ambiguity neutrality without the Principal Principle.

²⁸ Rowe and Voorhoeve (2018) similarly show—albeit by a very different line of argument—that there is a tension between Utilitarianism and ambiguity aversion.

itarian. More generally, if we allow for ambiguity sensitivity behind the veil, then the resulting distributive principle will not be what is sometimes called *Generalized Utilitarianism*, of which Prioritarianism and Utilitarianism are two special cases.²⁹ Generalized Utilitarianism holds that the preferences of a social planner should maximize:

Generalized Utility *There is a non-decreasing function π from $\mathcal{R}$ to $\mathcal{R}$ such that, for any $\mathcal{D} = \langle Ch(X_i) = \alpha_i \rangle$:*

$$GU(\mathcal{D}) = \sum_i \pi(u(X_i))\alpha_i \quad (8)$$

If ϕ is strictly increasing, then Generalized Utilitarianism is Prioritarianism if π is strictly concave, “Anti-Prioritarianism” if π is instead strictly convex, but Utilitarianism if π is linear. It is straightforward to verify that given the veil of ignorance setting, Desirability and the Principal Principle, Generalized Utilitarianism is equivalent to:

Generalized Chance Neutrality *There is a non-decreasing function π from $\mathcal{R}$ to $\mathcal{R}$ such that, for any gamble $\langle Ch(X_i) = \alpha_i \rangle$ and any X_i :*

$$u(X_i \cap \langle Ch(X_i) = \alpha_i \rangle) = \pi(u(X_i)) \quad (9)$$

Less formally, Generalized Chance Neutrality states that the utility of winning some prize X_i from a gamble $\langle Ch(X_i) = \alpha_i \rangle$ is an increasing function of the utility of X_i . But unlike Chance Neutrality, Generalized Chance Neutrality allows for the possibility that the utility of winning some prize X_i from a gamble $\langle Ch(X_i) = \alpha_i \rangle$ differs from the utility of X_i .

The equivalence between Generalized Utilitarianism and Generalized Chance Neutrality, given the above assumptions, has a peculiar³⁰ implication for Prioritarianism. This implication is informally explained below, but formally the implication is that Prioritarians must assume that for some X_i :

$$u(Ch(X_i) = 1) \neq u(X_i) \quad (10)$$

In other words, if the veil of ignorance argument is to support Prioritarianism, it has to be assumed (given Desirability and the Principal Principle) that the utility of a gamble that results in X_i for sure can be different from the utility of X_i . To take an example, the utility of a social gamble that will for sure result in you living at a particular level of welfare is, by 10, different from the utility of you living at that level of welfare. To take another example, the utility of a gamble that will for sure result in you gaining \$10 is, by 10, different from the utility of you gaining \$10.

The reason why Prioritarians must assume inequality 10 for some X_i is the following. Recall that according to standard Prioritarian thinking, the function π , which

²⁹ Note however that neither Distribution-Sensitive Utilitarianism nor Rank-Weighted Utilitarianism are special cases of Generalized Utilitarianism, since Generalized Utilitarianism is a generalization of *standard* Utilitarianism. See fn. 25.

³⁰ “Peculiar” might be an understatement, as the inequality 10 might be taken to show that the theory is not coherent (as Seth Lazar points out).

determines which weight to give to any level of welfare, is independent of which distribution is being evaluated; for instance, what determines the value that a life barely worth living contributes to the total value of a distribution is independent of whether everyone else in the distribution is better or worse off.³¹ Hence, the only way to prevent inequality 10 from ever holding, given Generalized Chance Neutrality, is to assume that in general, $\pi(u(X_i)) = u(X_i)$. For if this equality does not hold, then there will be some utility value for which inequality 10 holds.³² But if equality $\pi(u(X_i)) = u(X_i)$ does hold, then Generalized Chance Neutrality collapses into Chance Neutrality. And recall that by the veil of ignorance argument (and assuming Desirability and the Principal Principle), Chance Neutrality is equivalent to standard Utilitarianism. Hence, the only way to prevent Prioritarianism from collapsing into Utilitarianism, is to assume that inequality 10 holds for some X_i .

Before coming back to ambiguity sensitivity, I shall make two observations about the above inequality 10. First, the inequality 10 is the personal-gamble analogue of an obvious property of Prioritarianism, namely, that even when only one person is involved in say a policy intervention, Prioritarianism entails that the contributive value of that person's welfare to the moral value of the intervention might not be proportional to the difference that the intervention makes to the person's welfare. To put it differently, the Prioritarian value of a "distribution" of welfare over just one individual might not be proportional to (and hence, might not be equal to) that person's welfare.³³ As others have noted, this feature of Prioritarianism means that the theory may rank gambles that could only affect the welfare of a single individual differently from the individual's own (prudential) ranking of gambles.³⁴ Nevertheless, I suspect that many Prioritarians will find inequality 10 to be a rather difficult bullet to bite. After all, it would seem hard to accept that, say, the utility of a gamble that results in you gaining \$10 for sure is different from the utility of you gaining \$10.

Second, inequality 10 is not entailed by the failure of Chance Neutrality. To take a very simple example, Chance Neutrality is violated if for all α_i :³⁵

$$u(Ch(X_i) = \alpha_i) = u(X_i)\sqrt{\alpha_i} \quad (11)$$

³¹ This follows from an assumption of *separability in value across persons*, which is made by Prioritarians from Parfit (1991) to e.g. Holtug (2007) and Adler (2011). The typical Prioritarian justification for such separability is that the reason why benefits to the worse off matter more than benefits to the better off, is not that the former are worse off *than others*; rather, it has to do with their *absolute* level off well-(or ill-)being. As Holtug (2007, p. 132) puts it: "A benefit that falls at a particular level of welfare has the same moral value no matter what levels other individuals are at." (See also e.g. Parfit's 1991, p. 104 famous altitude analogy.)

³² To take an example, if $\pi(u(X_i)) = \sqrt{u(X_i)}$, then $u(Ch(X_i) = 1) \neq u(X_i)$ unless $u(X_i)$ equals 0 or 1.

³³ For instance, suppose $GU(\mathcal{D}) = \sum_i \sqrt{u(X_i)}\alpha_i$, which would make the evaluation of welfare distributions Prioritarian. Then even for a single person "distribution", where the person in question belongs to "group" X_i , $GU(\mathcal{D}) \neq u(X_i)$ unless $u(X_i)$ equals 0 or 1.

³⁴ See e.g. Rabinowicz (2002), Otsuka and Voorhoeve (2009), Voorhoeve and Fleurbaey (2012).

³⁵ To see that 11 violates Chance Neutrality, note that by Desirability and the Principal Principle, Chance Neutrality entails Linearity (Stefánsson and Bradley 2015), which states that the value of a gamble is a probability weighted sum of the utilities of the gamble's possible outcomes. Hence, Linearity entails that $u(Ch(X_i) = \alpha_i) = u(X_i)\alpha_i$, which rules out that Eq. 11 holds for all α_i .

But equation 11 is inconsistent with inequality 10. Hence, the failure of Chance Neutrality does not entail inequality 10.

Back then to the examination of the relationship between ambiguity sensitivity and theories of distributive ethics. Given what we have already learned about Chance Neutrality, it might not come as a surprise that Generalized Chance Neutrality is (given Desirability and the Principal Principle) inconsistent with ambiguity sensitivity. Therefore, if we assume, say, ambiguity aversion behind the veil of ignorance, then not only have we ruled out the possibility that the preferences that people display will support (standard) Utilitarianism; we have also ruled out that those preferences support Prioritarianism. For a straightforward corollary of Proposition 1 is that:

Proposition 2 *If a person satisfies Desirability, the Principal Principle, and Generalized Chance Neutrality, then she is ambiguity neutral*

For a concrete illustration of Propositions 1 and 2, and to see precisely what role (Generalized) Chance Neutrality plays, consider again the choice between bets on two coins, one of which is known to have an equal chance of resulting in heads as in tails, while the other is known to be either double headed or double tailed; where in each case one gets \$10 if the coin comes up heads but nothing otherwise. Let the bet on the fair coin coming up heads be $\mathcal{A} = Ch(\$10) = 0.5 \cap Ch(\$0) = 0.5$ and let the bet on the biased coin coming up heads be $\mathcal{B} = \{Ch(\$10) = 0 \cap Ch(\$0) = 1\} \cup \{Ch(\$10) = 1 \cap Ch(\$0) = 0\}$.

Finally, suppose that a person prefers $\mathcal{A}$ to $\mathcal{B}$ and satisfies Desirability and the Principal Principle. Furthermore, suppose that she is equally confident that the second coin is double headed as double tailed. Then according to her desirability function V , and simplifying the notation somewhat (see fn. 22), we have:

$$V(Ch(\$10 = 1) \cup Ch(\$10) = 0) < V(Ch(\$10) = 0.5) \quad (12)$$

So, by Desirability and the Principal Principle:

$$\begin{aligned} u(\$10 \cap Ch(\$10) = 1)0.5 + u(\$0 \cap Ch(\$10) = 0)0.5 \\ < u(\$10 \cap Ch(\$10) = 0.5)0.5 + u(\$0 \cap Ch(\$10) = 0.5)0.5 \end{aligned} \quad (13)$$

Hence, for this person to prefer $\mathcal{A}$ to $\mathcal{B}$, at least one of the following inequalities must hold:³⁶

$$u(\$10 \cap Ch(\$10) = 1) < u(\$10 \cap Ch(\$10) = 0.5) \quad (14)$$

$$u(\$0 \cap Ch(\$10) = 0) < u(\$0 \cap Ch(\$10) = 0.5) \quad (15)$$

But, according to (Generalized) Chance Neutrality, neither of the above two inequalities can hold, which illustrates why a person who satisfies Desirability and the Principal Principle will not be ambiguity averse if she satisfies (Generalized) Chance Neutrality.

³⁶ I assume that receiving \$10 dollars is preferred to receiving \$0. But that assumption is obviously not essential to the argument.

So, to summarize this section, we have found that if Stefánsson and Bradley's (2015, 2019) characterization of rational preference is correct, and by assuming general risk aversion, rational preferences behind the veil of ignorance support a form of Egalitarianism. We have also seen that the conditions necessary and sufficient for such preferences to support either Utilitarianism or Prioritarianism are inconsistent with people being ambiguity sensitive behind the veil. Hence, if we want to allow for ambiguity sensitivity, then we should neither take the veil of ignorance argument to support Utilitarianism nor Prioritarianism.

6 Concluding remarks

The findings of this paper raise the interesting question of *why* there is a connection between ambiguity aversion and Egalitarianism. The formal connection between risk aversion and Egalitarianism is, unlike the connection between ambiguity aversion and Egalitarianism, already well understood. Moreover, the connection between risk aversion and Egalitarianism is quite straightforward to intuitively understand, since the risk averse and the Egalitarian share the view that mean values do not completely determine the value of a distribution; the spread in values matters too. But once one recognizes that the ambiguity averse share this view too, as argued above, then the connection between ambiguity aversion and Egalitarianism should seem less mysterious. After all, for a given mean value, the ambiguity averse prefers the possible chance values to be as little spread—i.e., as *equal*—as possible.

Acknowledgements Open access funding provided by Stockholm University. This paper has been presented at Mini-Workshop on Attitudes to Risk at the Hebrew University of Jerusalem, Workshop on Ethics & Risk at the Australian National University, and the PPE seminar at the Institute for Futures Studies. I am grateful to the audience for their comments and questions. In addition, I am grateful to Richard Bradley, Philippe Mongin, and Alex Voorhoeve for stimulating and educational discussions about the content of this paper. Finally, many thanks to the special issue editors, Renee Bolinger and Seth Lazar, for helpful comments and a very enjoyable process, and also to two anonymous reviewers for very helpful suggestions for improving the paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix: General definitions and proofs

A.1. Definition of ambiguity sensitivity

Recall that an ambiguity averse person prefers as little spread as possible in the subjectively possible chances. So, we can for instance characterize ambiguity neutrality (i.e., neither aversion to nor love for ambiguity) as follows:³⁷

³⁷ As Machina and Siniscalchi (2014) note, there is no agreement on precisely what ambiguity sensitivity is. Hence, there is little hope for an agreed upon definition of ambiguity neutrality. The suggested definition

Definition (*Ambiguity neutrality*) First, I define two objective gambles $\mathcal{G}_\alpha$ and $\mathcal{G}_\beta$ over $\{X_1, \dots, X_n\}$:

- Let $\mathcal{G}_\alpha = \langle Ch_\alpha(X_i) = \alpha_i \rangle$,
- Let $\mathcal{G}_\beta = \langle Ch_\beta(X_i) = \beta_i \rangle$.

Second, I define a mixture of $\mathcal{G}_\alpha$ and $\mathcal{G}_\beta$, $\mathcal{G}_{\alpha\beta}$, for $p \in [0, 1]$, as follows:

- Let $\mathcal{G}_{\alpha\beta} = \langle Ch_{\alpha\beta}(X_i) = p\alpha_i + (1 - p)\beta_i \rangle$.

Third, I define a union of our objective gambles, which can be thought of as a subjective gamble between objective gambles:³⁸

- Let $\mathcal{G}_{\alpha\cup\beta} = \mathcal{G}_\alpha \cup \mathcal{G}_\beta$,
- where $\mathcal{G}_\alpha \cap \mathcal{G}_\beta = \perp$,
- and, according to our agent's subjective probability, P : $P(\mathcal{G}_\alpha \mid \mathcal{G}_{\alpha\cup\beta}) = p$, so $P(\mathcal{G}_\beta \mid \mathcal{G}_{\alpha\cup\beta}) = 1 - p$.

Finally, we say that the agent is ambiguity neutral, w.r.t. $\{X_1, \dots, X_n\}$, just in case for any p :

$$\mathcal{G}_{\alpha\beta} \sim \mathcal{G}_{\alpha\cup\beta} \quad (16)$$

To illustrate the above definition, consider again the coin example:

- Let $G_\alpha = \langle Ch_\alpha(\$10) = 1 \cap Ch_\alpha(\$0) = 0 \rangle$,
- Let $G_\beta = \langle Ch_\beta(\$10) = 0 \cap Ch_\beta(\$0) = 1 \rangle$,
- Let $G_{\alpha\beta} = \langle Ch_{\alpha\beta}(\$10) = p(1) + (1 - p)(0) \cap Ch_{\alpha\beta}(\$0) = p(0) + (1 - p)(1) \rangle$.

Now, if $p = 0.5$, then $G_{\alpha\beta}$ is equivalent to the bet on the coin that is known to be fair, and if $P(G_\alpha \mid G_\alpha \cup G_\beta) = P(G_\beta \mid G_\alpha \cup G_\beta)$, then $G_\alpha \cup G_\beta$ is the bet on the biased coin. So, the preference I have been discussing, $G_\alpha \cup G_\beta \prec G_{\alpha\beta}$, is not ambiguity neutral. Hence, we have evidence that the definition of ambiguity neutrality is apt.

Since ambiguity neutrality has not previously been defined for the present framework, it is worth comparing the above definition to Ellsberg's (1961) classical three-color example, where subjects know that 1/3 of the balls are red and that 2/3 are *either* black or yellow.

The frequently observed “Ellsberg preference” is: $g_2 \prec g_1, g_3 \prec g_4$ (Table 1).

For illustrative purposes, suppose that our subject believes that either all of the 2/3 are black or all of the 2/3 are yellow and is equally confident of each possibility. (The result below, i.e., that the Ellsberg preference is not ambiguity neutral given the above definition, does not however depend on this simplifying assumption.) Then given our framework (and ignoring the \$0 to keep things as simple as possible), we can formulate the gambles as:

is however in the spirit of an influential definition suggested by Schmeidler (1989). But it is modified in line with the framework favored in this paper (and it is a definition of ambiguity neutrality, rather than ambiguity aversion).

³⁸ Such unions corresponds in our framework to what is sometimes called “horse-roulette acts”, since they can be thought of as bets on a horse race where the prizes are to spin a roulette wheel, or “Anscombe–Aumann acts” (both terms referring to the seminal work of Anscombe and Aumann 1963).

Table 1 Ellsberg paradox

	Red	Black	Yellow
g_1	\$100	\$0	\$0
g_2	\$0	\$100	\$0
g_3	\$100	\$0	\$100
g_4	\$0	\$100	\$100

- $g_1 = \langle Ch_{g_1}(\$100) = 1/3 \rangle$
- $g_2 = \langle Ch_{g_2^1}(\$100) = 0 \rangle \cup \langle Ch_{g_2^2}(\$100) = 2/3 \rangle$
- $g_3 = \langle Ch_{g_3^1}(\$100) = 1/3 \rangle \cup \langle Ch_{g_3^2}(\$100) = 1 \rangle$
- $g_4 = \langle Ch_{g_4}(\$100) = 2/3 \rangle$

Now consider a mixture between the two “subjectively possible” chance distributions in g_2 :

$$Ch(\$100) = p(0) + (1 - p)(2/3) \quad (17)$$

When $p = 0.5$, $Ch(\$100) = p(0) + (1 - p)(2/3)$ is identical to g_1 . So, g_1 is a mixture between the two subjectively possible chance distributions in g_2 . Hence, a person with the Ellsberg preference $g_2 \prec g_1$ is not ambiguity neutral (given the above simplifying assumptions). An analogous argument shows $g_3 \prec g_4$ is not ambiguity neutral. Hence, we have further evidence that the proposed definition of ambiguity neutrality is apt.

A.2. Proof of main proposition

Proposition 1 *If a person satisfies Desirability, the Principal Principle, and Chance Neutrality, then she is ambiguity neutral.*

Proof Consider the gambles, $\mathcal{G}_\alpha \cup \mathcal{G}_\beta$, from the definition of ambiguity neutrality. By Desirability and the Principal Principle:

$$\begin{aligned} V(\mathcal{G}_\alpha \cup \mathcal{G}_\beta) &= p \left(\sum_i u(X_i \cap \mathcal{G}_\alpha) \alpha_i \right) + (1 - p) \left(\sum_i u(X_i \cap \mathcal{G}_\beta) \beta_i \right) \\ &= \sum_i u(X_i \cap \mathcal{G}_\alpha) (\alpha_i) p + \sum_i u(X_i \cap \mathcal{G}_\beta) (\beta_i) (1 - p) \end{aligned} \quad (18)$$

$$V(\mathcal{G}_{\alpha\beta}) = \sum_i u(X_i \cap \mathcal{G}_{\alpha\beta}) (p\alpha_i + (1 - p)\beta_i) \quad (19)$$

By Chance Neutrality, (18) and (19) become:

$$V(\mathcal{G}_\alpha \cup \mathcal{G}_\beta) = \sum_i u(X_i) (\alpha_i) p + \sum_i u(X_i) (\beta_i) (1 - p) \quad (20)$$

$$\begin{aligned}
 V(G_{\alpha\beta}) &= \sum_i u(X_i)(p\alpha_i + (1-p)\beta_i) \\
 &= \sum_i u(X_i)(\alpha_i)p + \sum_i u(X_i)(\beta_i)(1-p)
 \end{aligned} \quad (21)$$

Hence, since V represents $\precsim$, we have $G_\alpha \cup G_\beta \sim G_{\alpha\beta}$. $\square$

Observation. In footnote 27, I said that the importance of the Principal Principle for the result proven above is rather limited. To see this, note that had the mixture, in the definition of ambiguity neutrality, been defined in terms of the agent's subjective conditional probabilities rather than in terms of the objective probabilities, then we could have, without assuming the Principal Principle, proven that if a person satisfies Desirability and Chance Neutrality, then she is ambiguity neutral. (This can be seen by replacing the alphas and betas in the above proof with the relevant conditional subjective probabilities.) However, since I take it to be important for the concept of ambiguity that the mixture is an objective one, I find the suggested definition to be superior to one where the mixtures are defined in terms of subjective probabilities.

References

- Adler, M. D. (2011). *Well-being and fair distribution: Beyond cost-benefit analysis*. Oxford: Oxford University Press.
- Al-Najjar, N. I., & Weinstein, J. (2009). The ambiguity aversion literature: A critical assessment. *Economics and Philosophy*, 25(3), 249–284.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine. *Econometrica*, 21(4), 503–546.
- Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics*, 34(1), 199–205.
- Binmore, K. (2008). *Rational Decisions*. Princeton: Princeton University Press.
- Binmore, K., Voorhoeve, A., & Stewart, L. (2012). How much ambiguity aversion? *Journal of Risk and Uncertainty*, 45(3), 215–38.
- Bolker, E. D. (1966). Functions resembling quotients of measures. *Transactions of the American Mathematical Society*, 124(2), 292–312.
- Bradley, R. (2005). Bayesian utilitarianism and probability homogeneity. *Social Choice and Welfare*, 24(2), 221–251.
- Bradley, R. (2016). Ellsberg's paradox and the value of chances. *Economics and Philosophy*, 32(2), 231–248.
- Buchak, L. (2013). *Risk and rationality*. Oxford: Oxford University Press.
- Buchak, L. (2017). Taking risks behind the veil of ignorance. *Ethics*, 127(3), 610–644.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4), 643–669.
- Ellsworth, L. (1978). Decision-theoretic analysis of Rawls' original position. In A. Hooker, J. J. Leach, & E. F. McClennen (Eds.), *Foundations and applications of decision theory* (pp. 29–45). Dordrecht: D. Reidel.
- Fleurbaey, M. (2018). Welfare economics, risk and uncertainty. *Canadian Journal of Economics*, 51(1), 5–40.
- Gilboa, I., Postlewaite, A., & Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms? *Economics and Philosophy*, 25(3), 285–296.
- Harsanyi, J. C. (1953). Cardinal utility in welfare economics and in the theory of risk-taking. *Journal of Political Economy*, 61(5), 434–435.
- Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4), 309–321.

- Holtug, N. (2007). Prioritarianism. In N. Holtug & K. Lippert-Rasmussen (Eds.), *Egalitarianism: New essays on the nature and value of equality* (pp. 125–156). Oxford: Clarendon Press.
- Jeffrey, R. (1990/1965). *The logic of decision*. The University of Chicago Press.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in inductive logic and probability*. California: University of California Press.
- Machina, M., & Siniscalchi, M. (2014). Ambiguity and ambiguity aversion. *Handbook of the Economics of Risk and Uncertainty*, 1, 729–807.
- Mongin, P. (1995). Consistent Bayesian aggregation. *Journal of Economic Theory*, 66(2), 313–351.
- Mongin, P. (2001). The impartial observer theorem of social ethics. *Economics and Philosophy*, 17(2), 147–179.
- Mongin, P., & Pivato, M. (ms.). Social preference under twofold uncertainty.
- Otsuka, M., & Voorhoeve, A. (2009). Why it matters that some are worse off than others: An argument against the priority view. *Philosophy and Public Affairs*, 37(2), 171–199.
- Parfit, D. (2002/1991). Equality or priority? In Clayton, M., & Williams, A., (Eds.), *The ideal of equality*. New York: Palgrave.
- Rabinowicz, W. (2002). Prioritarianism for prospects. *Utilitas*, 14(1), 2–21.
- Rawls, J. (1999/1971). *A theory of justice*. Oxford University Press (revised edition).
- Rowe, T., & Voorhoeve, A. (2018). Egalitarianism under severe uncertainty. *Philosophy and Public Affairs*, 46(3), 239–268.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3), 571–587.
- Stefánsson, H. O., & Bradley, R. (2015). How valuable are chances? *Philosophy of Science*, 82(4), 602–625.
- Stefánsson, H. O., & Bradley, R. (2019). What is risk aversion? *The British Journal for the Philosophy of Science*, 70(1), 77–102.
- Trautmann, S. T., & van de Kuilen, G. (2015). *Ambiguity attitudes, chapter 3* (pp. 89–116). New York: Wiley.
- von Neumann, J., & Morgenstern, O. (2007/1944). *Games and economic behavior*. Princeton: Princeton University Press
- Voorhoeve, A., & Fleurbaey, M. (2012). Egalitarianism and the separateness of persons. *Utilitas*, 24(3), 381–398.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.