Population Ethics under Risk

Gustaf Arrhenius & H. Orri Stefánsson
Gustaf.Arrhenius@iffs.se
Orri.Stefansson@iffs.se

Version 2019-11-12

Abstract

Population axiology concerns how to evaluate populations in terms of their moral goodness, that is, how to order populations by the relations "is better than" and "is as good as". The task has been to find an adequate theory about the moral value of states of affairs where the number of people, the quality of their lives, and their identities may vary. So far, this field has largely ignored issues about uncertainty and the conditions that have been discussed mostly pertain to the ranking of risk-free outcomes. Most public policy choices, however, are decisions under uncertainty, including policy choices that affect the size of a population. Here, we shall address the question of how to rank population prospects—that is, alternatives that contain uncertainty as to which population they will bring about—by the relations "is better than" and "is as good as". We start by illustrating how well-known population axiologies can be extended to population prospect axiologies. And we show that new problems arise when extending population axiologies to prospects. In particular, traditional population axiologies lead to prospect-versions of the problems that they praised for avoiding in the risk-free settings. Finally, we identify an intuitive adequacy condition that, we contend, should be satisfied by any population prospect axiology, and show how given this condition, the impossibility theorems in population axiology can be extended to (non-trivial) impossibility theorems for population prospect axiology.

I. Introduction

Population axiology concerns how to evaluate populations in terms of their moral goodness, that is, how to order populations by the relations "is better than" and "is as good as". So far, discussions of population axiology have largely ignored the question of how to rank alternatives that contain uncertainty as to which population they will bring about. The adequacy conditions that have been discussed mostly pertain to the ranking of risk-free outcomes and are silent in regard to how to rank risky prospects. In other words, questions about how to rank what we call *population prospects*, that is, a population prospect axiology, have not yet been much addressed.[1]

This exclusion of uncertainty is unfortunate since most public policies that could affect the number and identities of future people contain a great deal of uncertainty as to which population they will bring about.[2] To take a concrete example, suppose we are considering increasing the cost of childcare, which has the predictable effect that some couples will decide to have a child just in case they happen to find well-paid jobs (and thus decreases the probability that these couples have a child). When evaluating this policy, we might want to consider the interests of the children who will come into existence just in case the people who will then be their parents find well-paid jobs. But how, if at all, should the probability that these children may, and may not, exist, affect our evaluation of the policy in question?

How to rank population prospects has important implications for public policy choices, as the above example illustrates.[3] For instance, if the probability of the non-existence of a possible person heavily detracts from the value of an alternative, then that gives us reason to resist any increase in the cost of childcare, but if the probability of the non-existence of a possible person is of, say, neutral value, then that makes it less objectionable to implement policies that are expected to have the effect that many possible and potentially very happy people will not come into existence.

Apart from the need to consider population prospects for policy purposes, there is a theoretical motivation for introducing uncertainty into population ethics. One might hope that removing the idealizing assumption of certainty could help us avoid the impossibility results that have been troubling population ethicists for decades (REFERENCES). After all, removing information can result in a situation where an adequacy condition that is used in a theorem no longer applies. So, removing information that is necessary for a condition to apply could be seen as one possible way in which one can avoid an impossibility result involving that condition.[4] In fact, that is what happens when one introduces risk and uncertainty into population ethics—that

[1] For important exception, see (Blackorby, Bossert, and Donaldson 2005),(Roberts 2007), (Asheim and Zuber 2016), (Nebel 2019), (Budolfson and Spears mimeo). (Voorhoeve and Fleurbaey 2016) and Nebel (2017) discuss prospects involving populations with different people but the same number of people. See also fn. 10.

[2] Contrary to the convention in economics, we use "uncertainty" and "risk" interchangeably throughout this paper. So, a situation of risk, and a situation of uncertainty, is one where one knows what outcome one's choices *might* result in, but does not know which outcome one's choice *will* result in. We shall assume that the probabilities that figure in our discussion are objective but nothing of substance hangs on this assumption.

[3] For further examples of this kind, see (Broome 1992, 2004, sec. 1.1., 2010, 2015). See also (Arrhenius forthcoming, ch 1) for a discussion.

[4] Cf. Arrhenius (forthcoming, sec. 6.4) discussion of how restriction on the measurement assumptions regarding welfare can help avoiding certain impossibility results in non-risky population axiology.

is, when one removes information about which state of the world is actual—since the traditional adequacy conditions for population axiology say nothing about how to rank what we call "non-trivial" prospects, that is, prospects that confer probabilities strictly between zero and one to different populations (as we further explain below).

An axiology for non-trivial prospect that is not plagued by impossibility results would moreover be a great victory from a practical point of view. Policies that could affect the size of a population typically contain considerable uncertainty, as previously mentioned. In fact, it is hard to think of *any* policy (population-affecting or not) that is certain to deliver a particular outcome. Therefore, one might think that, from a practical point of view, the impossibility results for risk-free population axiology are not so important. What is important, from a practical point of view, is the possibility of an adequate population prospect axiology. For if we have an adequate population prospect axiology, then we will be able to make consistent population policy choices (since these are always choices between prospects), notwithstanding the traditional impossibility results for risk-free populations.[5]

Unfortunately, it turns out that new problems appear when one tries to turn a population axiology into a population prospect axiology. Perhaps most surprisingly, we shall see that natural formulations of traditional (risk-free) population axiologies lead to prospect-versions of problems the avoidance of which in the risk-free setting is considered highly desirable features of these same axiologies.

Moreover, we shall identify an intuitive adequacy condition that, we contend, should be satisfied by any betterness ordering of population prospects, but which opens the door to impossibility results for population prospect axiology that are analogous to the traditional impossibility results for (risk-free) population axiology. Importantly, these results cannot be derived just from the (risk-free) adequacy conditions involved in the traditional impossibility theorems, and they hold for *non-trivial* prospects (as well as for trivial ones), i.e., for prospects that confer probabilities strictly between zero and one on different populations. So, contrary to the aforementioned hope, removing the idealizing assumption of certainty will unfortunately not help us avoid the notorious impossibility results in population ethics.

---

[5] Another reason for exploring what happens when we introduce probabilities in population ethics, is that we know that there are some multi-person optimisation problems where a solution can be found only if we introduce probabilities. In particular, there are some well-known problems in game theory where a *Nash equilibrium* exists only if the "players" are allowed to randomise between their available "pure" (i.e., risk-free) strategies. In a similar vein, one might hope that a solution can be found to the problems of population ethics by introducing probabilities.

II. Population Prospect Axiologies

A natural starting point in the search for a population prospect axiology is to consider some extensions of standard population axiologies to situations of uncertainty. But first, we need to introduce a few terms and distinctions used in traditional, risk-free, population axiology, which we will use when formulating the aforementioned extensions.

A *life*, as we shall be using the term, is individuated by, first, who leads the life, and, second, the welfare of the person. Hence, there is more than one possible life that a person could lead. We will assume that the welfare measure is normalised around the "neutral" life, such that lives worth living have positive welfare but lives worth not living have negative welfare. (But we shall remain agnostic as to what welfare is and how to measure it.) *Populations* are individuated by the lives they contain. Therefore, populations are individuated by, first, the people they contain, and, second, the welfare of these people. Hence, two different populations could be made up of exactly the same people.

In traditional population axiology, one distinguishes between the welfare of a life, i.e., how good or bad it is for the person living it, and the contributive value of a life. By the contributive value of a life we shall mean the value that a life confers on a population of which it is a member. More exactly, the contributive value of a life x relative to a population A, of which x is a member, is the difference in value between A and the population consisting of all the A-lives except x. This distinction is useful for population prospect axiology too, for we can now reformulate one of our main questions as "What is the contributive value of a *possible* life to a prospect?"

With this terminology in hand, let's consider how the traditional population axiologies could be extended to situations of uncertainty. We start with Total Utilitarianism, according to which the contributive value of a person's life equals her welfare, and the value of a population is calculated by summing the welfare of all lives in the population.

$$TU(X) = \sum_{i=1}^{n} u_i = u_1 + u_2 + \ldots + u_n$$

In the above formula, n is the size of population X and u_i is a numerical representation of the welfare of the ith life in population X.

A natural way to turn Total Utilitarianism into a population prospect axiology, is to sum up weighted Total Utilitarian values of all the possible populations that a prospect might result in, where the weight on the value of each population is determined by the probability that the prospect will result in that population.[6] Let's call the resulting axiology Prospect Total Utilitarianism.

$$PTU(P) = \sum_{i=1}^{m} p(s_i)TU\big(P(s_i)\big) = p(s_1)TU(P(s_1)) + p(s_2)TU(P(s_2)) + \cdots + p(s_m)TU(P(s_m))$$

In the above formula, the s_i are mutually incompatible states of the world that determine in which population a population prospect results;[7] $P(s_i)$ is the population resulting from population prospect P in state s_i; and $p(s_i)$ is the probability of s_i. So, according to Prospect Total Utilitarianism, the contributive value of a possible life to a prospect is positive (negative, neutral) if the life would have positive (negative, neutral) welfare and proportional to the probability of its realisation by the prospect.

One of the great challenges that examinations of different size populations have raised for Total Utilitarianism, is how to respond to the fact that the latter gives rise to what Parfit (1984) called the *Repugnant Conclusion*:

> *The Repugnant Conclusion*: For any perfectly equal population consisting of people with very high positive welfare, there is a population consisting of people with very low positive welfare which is better, other things being equal.[8]

---

[6] The population prospects axiologies we formulate treat risk and uncertainty as standard expected utility theory does. There are reasons for being sceptical of that treatment; for instance, since it does not allow for risk aversion with respect to utility. However, to keep things simple, we shall assume the standard expected utility treatment of risk. It is easy to see that the results we shall discuss would still hold if we assumed some of the *minimally normative* alternative treatments of risk that have recently been suggested, that is, treatments that satisfy transitivity and stochastic dominance, such as the theories suggested by Quiggin (1982), Kahneman and Tversky (1992), Buchak (2013) and Stefánsson and Bradley (2015, 2019). In contrast, our results do not hold if one uses, say, the original *Prospect Theory* (Kahneman and Tversky 1979), which cannot simultanously satisfy stochastic dominance and transitivity, or *Regret Theory* (Bell 1982, Loomes and Sugden 1982), which only satisfies transitivity w.r.t. a triple of alternatives if at least one alternative in the triple strictly dominates another alternative in the triple (Diecidue and Somasundaram 2017). While the latter theories may be accurate descriptions of how people actually order prospects, few if any scholars would suggest them as theories about how one ideally *should* order prospects. Therefore, we think that we can justifiably set them aside for the present purposes, where the aim is to discuss how one should order population prospects.

[7] In other words, states are pairwise incompatible but collectively exhaustive contingencies.

[8] See (Parfit 1984, p. 388). This formulation is more general than Parfit's except that he doesn't demand that the people with very high welfare are equally well off. The *ceteris paribus* clause means that the compared populations are roughly equal in all other putatively axiologically relevant aspect apart from individual welfare levels. Hence, other values and considerations are not decisive for the value comparison of involved populations. Although it is through Parfit's writings that this implication of Total Utilitarianism became widely discussed, it was already noted by Henry Sidgwick (1907), p. 415, before the turn of the century (he didn't claim that it was "repugnant", however). For other early sources of the Repugnant Conclusion, see (McTaggart 1927, pp. 452-3; Broad 1930, pp. 249-50; and Narveson 1967).

It should be evident that the prospect version of Total Utilitarianism also leads to the Repugnant Conclusion not only in non-risky situations but also in risky ones. For instance, for any prospect *P*, all of whose possible outcomes are populations with lives with very high welfare, Prospect Total Utilitarianism will recommend choosing instead some prospect *Q*, all of whose possible outcomes are larger populations with lives with very low positive welfare. Here, the actual outcome, if we follow the recommendation of Prospect Total Utilitarianism, will be an instance of the Repugnant Conclusion. Likewise, for other undesirable implications of Total Utilitarianism in non-risky settings, such as the Very Repugnant Conclusion.[9]

Moreover, Prospect Total Utilitarianism violates probabilistic analogues of the traditional (risky-free) *Quality* condition which essentially rules out the Repugnant Conclusion (and which we state in section 4). Consider the following probabilistic analogue, which can be derived from the non-probabilistic condition with help of a *dominance* condition that we discuss in section 4:

> *Probabilistic Quality*: There is a perfectly equal population with very high welfare, *A*, such that for some population *X* and any probability *p* and any population *C* with very low welfare (positive or negative), a prospect that has a probability of *p* of resulting in *A* and a probability of 1-*p* of resulting in *X*, is better than a prospect that has a probability *p* of resulting in *C* and probability 1-*p* of resulting in *X*.

We could weaken the condition by for instance requiring *p* to be arbitrarily close to 1. But whether we weaken it or not, Prospect Total Utilitarianism violates the condition. And the reason is the same as the reason why Total Utilitarianism leads to the Repugnant Conclusion, namely, that if a population with low average welfare is sufficiently large, then it is by what we might call "total views" deemed better than a smaller population consisting of people with very high welfare. Therefore, if *C* in the statement of *Probabilistic Quality* is sufficiently large, then even though people in *C* have low welfare, any prospect between C and *X* will be deemed better than a prospect between A and *X*, irrespective of *X*.

In comparison with other axiologies, Total Utilitarianism handles (non-risky) cases involving negative welfare better than most alternative axiologies, such as Average and Critical-Level Utilitarianism (more on which below). For example, Total Utilitarianism clearly avoids the *Very Sadistic Conclusion* which is implied by Critical-Level Utilitarianism:[11]

---

[9] *The Very Repugnant Conclusion:* For any perfectly equal population with very high positive welfare, and for any number of lives with very negative welfare, there is a population consisting of the lives with negative welfare and lives with very low positive welfare which is better than the high welfare population, other things being equal.

[11] For a discussion, see (Arrhenius forthcoming, 2000a, 2000b). Average Utilitarianism has similar problems.

The Very Sadistic Conclusion: For any population with negative welfare, there is a population with positive welfare which is worse, other things being equal

Unlike Total Utilitarianism, Prospect Total Utilitarianism has counterintuitive implications when it comes to people with negative welfare. For instance, the view leads to what we might call the *Risky Very Sadistic Conclusion*:[12]

The Risky Very Sadistic Conclusion: For any arbitrarily high probability $p < 1$ and any population A consisting of lives with very high positive welfare, and any population B consisting of lives with very negative welfare, there is a population C consisting of lives with very low positive welfare such that a prospect that with probability p results in B and with probability 1-p results in C is better than a prospect that will for sure result in A.

The following diagram illustrates the above result:

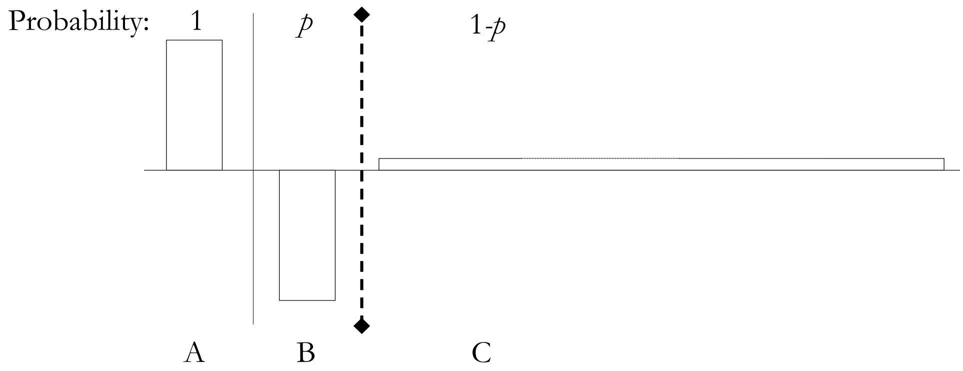

Figure 1.

In this diagram, the A-people have very high positive welfare, the B-people have very negative welfare, and the C-people have very low positive welfare. (The width of the blocks represents the number of lives in the population, the height represents their welfare; dashes indicate that the block in question should strictly speaking be much wider than shown.) The diagram

[12] We note that Budolfson and Spears (mimeo) discuss a more general result, which they call the *Probabilistic Very Repugnant Conclusion*. A proper comparison of the results we discuss here to the results discussed by Budolfson and Spears will have to await another occasion.

represents two prospects, indicated by the unbroken vertical line, one prospect that results in A for sure and another that results in either B or C (the dashed line separates possible outcomes in the latter prospect). Now, for any probability p, however high, if there are sufficiently many C-people, then the expected total welfare of prospect P, which has probability p of resulting in B and probability 1-p of resulting in C, will be higher than the expected total welfare of prospect Q, which results in A for sure. Thus, Prospect Total Utilitarianism deems P better than Q. This holds irrespective of how much the B-people suffer and of how many they are. In other words, while Total Utilitarianism avoids the Very Sadistic Conclusion, *Prospect* Total Utilitarianism recommends prospects that will almost certainly result in a "very sadistic conclusion" (hence the name of the conclusion).[13]

Now, one might think that the above problem is just a general implication of *Expected Utility Theory* (EUT), that has nothing specifically to do with population ethics. After all, EUT, which inspired the above way of turning Total Utilitarianism into a prospect axiology, implies that for any desirable outcome X, there is a fantastic outcome Y and a terrible outcome Z, such that for any probability $p<1$, one should prefer a prospect that has probability p of resulting in Z and probability 1-p of resulting in Y, to a "prospect" that is sure to result in X. In other words, EUT has implications with the same structure as the Risky Very Sadistic Conclusion.

In response, we would like to point out that while one might see the Risky Very Sadistic Conclusion as a special case of a more general implication of EUT, that does not change the fact that the special case in question is illuminating for debates in population ethics. In particular, it teaches us something important about Total Utilitarianism, namely, that while the theory may deliver the intuitively right result for some paradigmatic cases involving negative welfare under certainty, it leads to counterintuitive results for cases with negative welfare under risk.[14]

Faced with the Repugnant Conclusion and other similar problems, a Utilitarian might be tempted to opt for a non-total version of the theory, such as Average Utilitarianism or Critical-

---

[13] Some may dislike the name of this conclusion for the following reason: Total utilitarians think that the population with negative welfare is the *worst* of the possible populations involved in the above prospect; hence, the Risky Very Sadistic Conclusion differs in an important respect from the ordinary Very Sadistic Conclusion, where critical-level utilitarians think that a population with negative welfare is *better* than one with positive welfare. (Thanks to […] for bringing this objection to our attention.) Nevertheless, we think that the label is apt for the reason explained above: The Very Risky Sadistic Conclusion will almost certainly result in an outcome analogous to the outcome of the Very Sadistic Conclusion.

[14] One might argue that this means that while Total Utilitarians have traditionally bashed Average Utilitarians with the Very Sadistic Conclusions, Average Utilitarians now have a simple response: The most natural extension of Total Utilitarianism to situations of risk and uncertainty will endorse prospects that have a probability arbitrarily close to 1 of resulting in the very same outcome as the "sadistic" one in the Very Sadistic Conclusion. There are of course possible responses that Total Utilitarians could make (e.g. opting for a factualist criterion of betterness and leaving prospects problems to their preferred decision method). We shall not pursue this issue further here. For a discussion of the distinction between criterions and decision methods, see (Bales 1971), (Brink 1986 pp. 421–7), and (Danielsson 1974 pp. 28–9).

Level Utilitarianism, which can respectively be extended to Prospect Average Utilitarianism and Prospect Critical Level Utilitarianism. It is easy to see (e.g. below) that these theories avoid the Repugnant Conclusion. However, as is well known, these alternatives to Total Utilitarianism have, in the risk-free setting, their own counterintuitive conclusions to deal with; conclusions that are arguably at least as counterintuitive as the Repugnant Conclusion.[15] It should be evident that since Prospect Average Utilitarianism and Prospect Critical Level Utilitarianism are extensions of Average Utilitarianism and Critical Level Utilitarianism to risky situations, they inherit in risk-free situations the problems of their ancestors.[16] Moreover, it turns out that Prospect Average Utilitarianism and Prospect Critical Level Utilitarianism—and, in fact, any population prospect axiology that satisfies transitivity and a weak dominance condition—faces additional problems due to the introduction of uncertainty, as we shall see in section IV.

To appreciate some of the problems that non-total Utilitarian theories are faced with, let's look more closely at Average Utilitarianism, according to which the value of a population is found by summing up the welfare of the lives in the population and dividing this sum with the number of lives in the population.[17] (After examining Average Utilitarianism and its prospect extension, we shall briefly discuss Critical Level Utilitarianism and its prospect extension.) Hence, the contributive value of a life depends on the average welfare of the population of which it is a member: The contributive value of a life is positive (negative, neutral) if the life has a higher (lower, same) welfare than the average welfare in the population.

$$AU(X) = \sum_{i=1}^{n} u_i \times \frac{1}{n}$$

Now, we can extend this theory to situations of risk—thus creating *Prospect Average Utilitarianism*—in an analogues way to how Total Utilitarianism was extended to Prospect Total Utilitarianism:

$$PAU(P) = \sum_{i=1}^{m} p(s_i) AU\big(P(s_i)\big)$$

[15] See, e.g., (Parfit 1984; Arrhenius 2000a, forthcoming).

[16] The same holds for the recent prospect axiology presented in (Asheim and Zuber 2016) which is an extension of the population axiology in (Asheim and Zuber 2014). The latter violates some arguably compelling adequacy conditions, such as the Non-Extreme Priority Condition (Arrhenius forthcoming).

[17] The problem with Average Utilitarianism that we focus on here also rises for Critical Level Utilitarianism. See (Arrhenius 2000b, 2000a, forthcoming).

Average Utilitarianism gives rise to the *Sadistic Conclusion*:[18]

> *The Sadistic Conclusion*: It can be better to add to a population people with negative rather than positive welfare, other things being equal.

After all, adding a few lives with negative welfare to population X could affect the average welfare of the resulting population less negatively than adding to X a lot of lives with positive welfare that is nevertheless lower than the average welfare in X. Even worse, Average Utilitarianism violates the following seemingly unassailable principle, which is even logically weaker than what is required to avoid the Sadistic Conclusion:[19]

> *Weak Non-Sadism*: There is a negative welfare level and a number of lives at this level such that an addition of any number of people with positive welfare is at least as good as an addition of the negative welfare lives, other things being equal.

Prospect Average Utilitarianism has the above counterintuitive consequences as well, most straightforwardly in the special case where one population has probability 1. Furthermore, Prospect Average Utilitarianism violates the following probabilistic analogue of Weak Non-Sadism, even when the probabilities involved are strictly between 0 and 1:

> *Probabilistic Weak Non-Sadism*: For any background population, for any population X, and for any population Y consisting of lives with very positive welfare, and for any probability p, there is a very negative welfare level and a number of lives at this level, call this population Z, such that a prospect that has probability p of resulting in the addition of X to the background population and probability 1-p of resulting instead in the addition of Y, is better than a prospect that has probability p of resulting in the addition of X to the background population and a probability 1-p of resulting instead in the addition of Z.

The reason why Prospect Average Utilitarianism violates Probabilistic Weak Non-Sadism is the same as the reason why Average Utilitarianism violates Weak Non-Sadism. Take any population

18 See, e.g., (Arrhenius 2000a).

19 A welfare level is simply an equivalence class on the set of possible lives.

A where, say, the majority has very high welfare. Then according to what we might call "average views", if a subpopulation of *A* consisting of people with negative welfare is sufficiently small, then its effect on the overall value of population *A* is less negative than the effect of a larger subpopulation of *A* with, say, only slightly positive welfare. So, in particular, if the population *Z* in the statement of Probabilistic Weak Non-Sadism is sufficiently small, then although the people in *Z* have negative welfare, a Prospect Average Utilitarian should prefer risking adding *Z* to any background population to risking adding *Y* to the background population.

Moreover, while Average Utilitarianism (and Critical Level Utilitarianism) of course avoids the risk-free Repugnant Conclusion, Prospect Average Utilitarianism (and Prospect Critical Level Utilitarianism as well as of Prospect Total Utilitarianism) leads to what we might call the *Risky Repugnant Conclusion*:

> *The Risky Repugnant Conclusion*: For any arbitrarily low probability $p>0$, any perfectly equal population A consisting of people with very high positive welfare, and any population C consisting of people with low welfare (including both slightly positive and very negative welfare), there is a population B consisting of lives with very high welfare, higher than in A, such that a prospect that with probability p results in B and with probability 1-p results in C is better than a prospect that will for sure result in A.

The following diagram illustrates the above result:

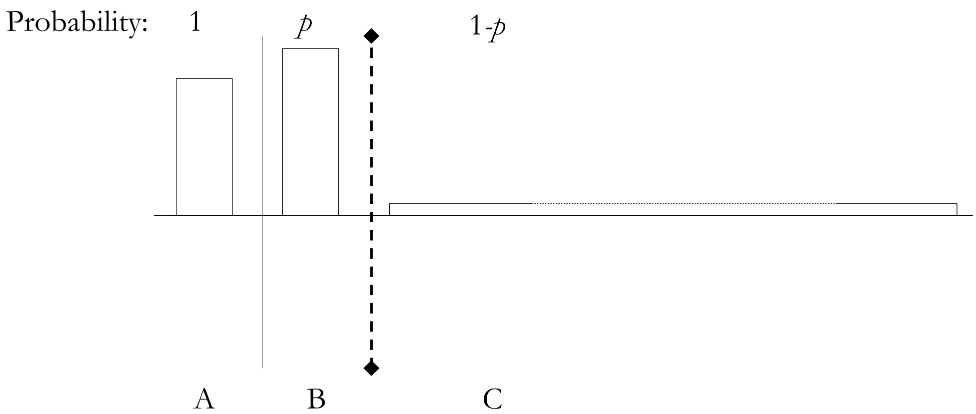

Figure 2.

The above result means that while Average Utilitarianism (and Critical Level Utilitarianism) avoids the Repugnant Conclusion, *Prospect* Average Utilitarianism (and Prospect Critical Level

Utilitarianism) recommends prospects that will almost certainly result in a Repugnant Conclusion, that is, in the special case where C is a huge population consisting of lives with very low positive welfare.

For the sake of completeness, and to convince the reader that Prospect Critical Level Utilitarianism really does lead to the Risky Repugnant Conclusion, it is worth formally stating the theory. The general intuition behind Critical Level Utilitarianism is that only people with welfare that is above some "critical level", which is typically assumed to be higher than the point at which life becomes worth living, can increase the value of a population. More precisely, the risk-free version of the theory states that for some critical level $c>0$,[20] the value of a population X is given by:

$$CLU(X) = \sum_{i=1}^{n} u_i - c$$

Blackorby et.al. (2005) have extended standard Critical Level Utilitarianism to a prospect theory by essentially applying expected utility theory to prospects involving populations where each population is valuated by the above formula. That is, they extend Critical Level Utilitarianism in the same way as we have extended Average and Total Utilitarianism (albeit using a somewhat different notation and formal framework than we do). The resulting theory states that:

$$PCLU(P) = \sum_{i=1}^{m} p(s_i) CLU\big(P(s_i)\big)$$

It should be evident that while Critical Level Utilitarianism was partly designed to avoid the Repugnant Conclusion, Prospect Critical Level Utilitarianism leads to the Risky Repugnant Conclusion. For as long as the people in the B population are sufficiently high above the critical level, then Prospect Critical Level Utilitarianism will instruct us to gamble on B rather than getting A for sure, irrespective of how much more likely C is than B.

Now, Average or Critical Level Utilitarians could avoid this result by postulating that there is a limit to how high a person's welfare can be. But what we take this and the previous result to illustrate—and the point we want to emphasise—is that new problems appear when we try to turn a population axiology into a population prospect axiology. For, as we have seen, natural

[20] It is worth noting that c need not be a constant, but can be a function of the size of the evaluated population (Blackorby et.al. 2005).

formulations of each of the three standard versions of Utilitarianism lead to prospect-versions of the conclusions that they avoid in the risk-free settings.[21]

For the remainder of this paper, the question we will focus on is whether the impossibility results that have been proven for risk-free population axiology—i.e., results showing that no population axiology can satisfy different sets of intuitively compelling adequacy conditions—also hold in risky settings. As has been shown in risk-free population axiology, there is a great advantage to focusing on exploring the logical consistency of different adequacy conditions, rather than formulating new theories in response to the deficiencies that have been found in the old theories. One can always formulate a theory that avoids a certain counterintuitive conclusion, but when it comes to population ethics, we have results showing that such a theory will imply some other unappealing result (some of which we discuss below). Hence, we risk chasing an illusory goal forever. With the adequacy-condition-approach, however, we will know with certainty which sets of considered intuitions are collectively consistent.

In the spirit of our favoured approach, we shall, in the next section, introduce a weak and very plausible condition that ensures that the ranking of population prospects coheres with the ranking of populations. As we show in section IV, this condition suffices to generate impossibility results for population prospect axiology that are similar to the traditional impossibility results for (risk-free) population axiology. Moreover, as we shall explain in section IV, something like the condition introduced in the next section is *needed* to generate impossibility results for population *prospect* axiologies, since the traditional adequacy conditions for risk-free population axiology say nothing whatsoever about how to rank prospects.

III. Population State Dominance

As an illustration of a rather straightforward comparison of population prospects, consider prospects P_1 and P_2 in Table 1. We shall assume, both in the below example and throughout this paper, that the only axiologically relevant features at stake are, first, the welfare of the people in the populations that the prospects might result in, and, second, the probabilities that the prospects result in these populations.[22] Moreover, in the example below, we assume that the only person whose welfare P_1 and P_2 affect is Ann, who will for sure exist and lead a horrible and suffering life (denoted by "Ann: -100"), which we assume is not worth living, if we choose P_2. The same is true

[21] However, note that even in the non-risky case, Average Utilitarianism implies repugnant-like conclusions. See e.g., (Arrhenius forthcoming, 2000b).
[22] Cf. fn. 8.

if we choose P_1 and state of the world s_1 obtains. However, in state s_2 the choice of P_1 means that Ann will not exist (which we denote by 'Ω').

	s_1	s_2
P_1	Ann: -100	Ω
P_2	Ann: -100	Ann: -100

Table 1.

It seems evident that P_1 is better than P_2. After all, Ann is sure to lead a suffering life given the choice of P_2, but she could be spared these sufferings given the choice of P_1.

The judgement that P_1 is better than P_2 follows from the combination of two very compelling principles. One of these principles, *Negative Mere Addition*, is an uncontroversial and generally accepted adequacy condition for (risk-free) population axiologies, and should also be uncontroversial, we surmise, for population prospect axiologies:[23]

> *Negative Mere Addition*: An addition of lives with negative welfare makes a population worse, other things being equal.

The other principle is an instance of the well-known *State-Wise Dominance* principle from decision theory. The role of this principle, which is rarely questioned as a normative constraint, is to ensure that the ranking of prospects coheres with the ranking of the outcomes (in our case, the populations) that the prospects could result in.[24] We can see no reason why one would reject this as an adequacy condition for population prospect axiologies, assuming, as we have done, that the only axiologically relevant feature at stake is the welfare of the people involved and the probabilities of the possible populations (more on this below).[25]

Now consider again the two prospects in Table 1, P_1 and P_2. Since we individuate populations according to the people in them and the people's welfare, prospects P_1 and P_2 result in the same population if state s_1 obtains. That is, given the previous notation, $P_1(s_1)=P_2(s_1)$. And, of course,

---

23 Recall that we stipulate that lives with negative welfare are worth not living.

24 The classical theories of rational choice under uncertainty—such as Ramsey's (1926), Savage's (1954), and Jeffrey's (1965)—all impose State-Wise Dominance for the type of situations we are considering, i.e., situations where the states' probabilities are unaffected by the agents' choice.

25 It might be worth mentioning that unlike *Event*-Wise Dominance, State-Wise Dominance does not rule out intuitively reasonable risk attitudes that that are inconsistent with standard expected utility theory, such as the so-called *Allais Preference* (named after Maurice Allais 1953).

every population is as good as itself. But even if this were not how we wanted to individuate populations, we would still find that $P_1(s_1)$ and $P_2(s_1)$ are equally good, assuming that we accept the following very plausible principle: if populations X and Y consist of the same people and each person has the same welfare in X as in Y, then X and Y are equally good (other things being equal). So, $P_1(s_1)$ and $P_2(s_1)$ are equally good. However, $P_1(s_2)$ is better than $P_2(s_2)$ according to Negative Mere Addition. So, if s_1 obtains then P_1 and P_2 are equally good but if s_2 obtains then P_1 is better than P_2. Hence, we have a dominance argument in favour of P_1 over P_2.

The principle that underwrites the above dominance reasoning is, as previously mentioned, an instance of the traditional State-Wise Dominance principle. Let's call this instance of the principle *Population State Dominance*:

> *Population State Dominance*: If for all states s_i, population $P_1(s_i)$ is at least as good as $P_2(s_i)$, then prospect P_1 is at least as good as P_2, other things being equal; and if, in addition, for at least one non-null[26] state s_j, population $P_1(s_j)$ is better than $P_2(s_j)$, then prospect P_1 is better than P_2, other things being equal.

It might be worth noting that just like State-Wise Dominance, Population State Dominance entails that the value of an outcome cannot depend on what other outcomes the prospect might result in, assuming that the outcome has been described in in sufficient detail for the application of the principle to be legitimate—in our case, assuming that the value of the relevant population has been described in sufficient detail. In other words, the principle rules out value dependencies across different states of the world.[27] We can think of only one (ultimately misguided) objection to this constraint in the present context. Suppose populations A and B consist of different but equally many people, all having the same very high welfare, and that we deem A and B to be equally good in all other respects. And imagine that we are comparing a prospect that will result in A for sure (i.e., the prospect results in A in all states of the world) with a prospect that that has, say, a 0.5 chance of resulting in A and a 0.5 chance of resulting in B (i.e., the prospect results in A in some states of the world but results in B in other states). Then many people have the intuition that the second prospect is strictly better than the first, which might seem to violate Population State

[26] A state is non-null if it has some chance of obtaining.

[27] Note however that the fact that knowledge of what could have been can affect people's welfare is not an argument against Population State Dominance, since such effects would be reflected in the individual welfare measure.

Dominance. After all, the second prospect but not the first gives the B-people a chance to exist and lead lives worth living.[28]

Should we give up Population State Dominance in response to the above objection? We think not. Note that the principle, as we have stated it, has a *ceteris paribus* clause: "other things being equal". Now, if there is value to giving a person a chance of existing with positive welfare that is not exhausted by what is good for the person in question, then other things are *not* equal in the above purported counterexample to Population State Dominance. Hence, Population State Dominance would *not* say that the prospect that results in A for sure is at least as good as the prospect that has an equal chance of resulting in A as in B. So, the value of giving people a chance does not undermine Population State Dominance. For the purpose of our argument, however, we can assume that things other than the welfare, number, and identities of individuals are equal. In particular, the importance of the results we shall discuss in the next section is not threatened by this assumption.

Perhaps additional counterexamples to Population State Dominance will come to light. In that event, we would have three options: First, to examine whether the examples fall under the *ceteris paribus* clause; second, to examine whether the purported counterexamples can be avoided by building the value dependencies into the individual welfare measure;[29] and, finally, if the first two attempts to rescue Population State Dominance fail, we could examine whether the principle can be plausibly weakened. However, since we are not aware of any further objections to Population State Dominance, we will continue to take it for granted.

Let's return then to our example from Table 1. Note that since "is at least as good as" is a reflexive relation, any population is at least as good as itself. Hence, Population State Dominance and Negative Mere Addition together entail the following principle, which we have been implicitly assuming:[30,31]

---

[28] Diamond (1967) uses an example with this structure to criticise the much stronger *Sure Thing Principle*. For further discussion, see e.g. Broome (1991) and Stefánsson (2015).

[29] Stefánsson (2015) shows how one can formally construct such a measure.

[30] More generally, Population State Dominance and Negative Mere Addition together entail:

> *General Negative Existential Risk Dominance*: If for all states s_i such that i≠j, $P_1(s_i)$ is equally as good as $P_2(s_i)$, and there is a state s_j such that the only difference between $P_1(s_j)$ and $P_2(s_j)$ is that $P_2(s_j)$ contains a life with negative welfare that $P_1(s_j)$ does not contain, then P_1 is better than P_2, other things being equal.

[31] It is worth pointing out that the difference between P_1 and P_2, i.e., a life with negative welfare and positive welfare, is outside of the set of factors which are assumed to be equal across the prospects we consider. For as the reader may recall, that set does not include the number and identities of people. Therefore, Population State Dominance does entail, in the presence of Negative Mere Addition, that P_1 is better than P_2.

Negative Existential Risk Dominance: If for all states s_i such that i≠j, $P_1(s_i)$ is identical to $P_2(s_i)$, and there is a state s_j such that the only difference between $P_1(s_j)$ and $P_2(s_j)$ is that $P_2(s_j)$ contains a life with negative welfare that $P_1(s_j)$ does not contain, then P_1 is better than P_2, other things being equal.

Since both Negative Mere Addition and Population State Dominance should be taken as adequacy conditions for a satisfactory population prospect axiology, we contend, any adequate population prospect axiology should satisfy Negative Existential Risk Dominance, and should thus deem P_1 to be better than P_2.

How do the population prospect axiologies discussed above fare with respect to Negative Mere Addition and Population State Dominance? Population State Dominance is satisfied by all the above-mentioned theories, and, we contend, by any plausible prospect axiology (with the caveats discussed above). Prospect Total Utilitarianism and Prospect Critical Level Utilitarianism also satisfy Negative Mere Addition and thus satisfy Negative Existential Risk Dominance. However, Prospect Average Utilitarianism violates Negative Mere Addition (just like Average Utilitarianism does), and, indeed, Negative Existential Risk Dominance.

As an illustration of why Prospect Average Utilitarianism violates Negative Existential Risk Dominance, consider again prospects P_1 and P_2 in Table 1, and recall that the only difference between $P_1(s_2)$ and $P_2(s_2)$ is that the latter contains Ann leading a horrible and suffering life—with welfare -100—whereas Ann does not exist in the former (no other person is affected by the choice between these two prospects). Now, if Ann's welfare in $P_2(s_2)$, although negative, is higher than the average welfare in $P_2(s_2)$ when Ann's welfare has been excluded, then the addition of Ann with this negative welfare improves the average welfare of the population. Hence, $P_2(s_2)$ is better than $P_1(s_2)$, according to Average Utilitarianism.[32] Finally, since according to Prospect Average Utilitarianism, the value of P_1 is a (possibly weighted)[33] average between the values of $P_1(s_1)$ and $P_1(s_2)$, and similarly for P_2, Prospect Average Utilitarianism deems P_2 better than P_1, in violation of Negative Existential Risk Dominance.

[32] One might think that the Average Utilitarian value of $P_1(s_2)$ is undefined since nobody exists in s_2. That would mean that $P_1(s_2)$ could not be better than $P_2(s_2)$ according to Average Utilitarianism. However, that would be to misunderstand the example. We are not assuming that *only* Ann can exist, and thus that *nobody* exists in $P_1(s_2)$. Rather, we are assuming that only Ann is affected by the choice between P_1 and P_2.

[33] That is, if s_1 and s_2 are not equiprobable.

IV. A Population Prospect Impossibility Theorem

To answer the general question about the possibility of an adequate population prospect axiology, we can exploit the generally accepted and much studied adequacy conditions for (risk-free) population axiology, and add to them Population State Dominance, which, as previously mentioned, we take to be an adequacy condition for population prospect axiology. Although the traditional (risk-free) adequacy conditions are known to be mutually inconsistent as conditions on how to rank outcomes, they are not mutually inconsistent as conditions on how to rank prospects, since none of them says anything about how to rank prospects. Hence, a condition like Population State Dominance is needed, to connect the ranking of outcomes to the ranking of prospects.

We shall focus on three adequacy conditions for (risk-free) population axiology that generate a very simple and easy to demonstrate impossibility theorem, which we will extend to population prospect axiology. Some of these conditions can be and have been questioned. However, as will become evident, the various other impossibility theorems that have been proven with logically weaker and intuitively more compelling adequacy conditions for population axiology (see, e.g., Arrhenius forthcoming) can be extended to population prospect axiology in a way analogous to how we extend the simple theorem. So, the aim here is not to prove an impossibility theorem involving only intuitively compelling adequacy conditions, but to illustrate how one can translate all theorems for (risk-free) population axiology into theorems for prospect population axiology on a domain that contains non-trivial prospects (i.e., prospects that confer chances strictly between 0 and 1 on different populations).

The three adequacy conditions are:

> *Egalitarian Dominance*: If A is a perfectly equal population of the same size as population B, and every person in A has higher welfare than every person in B, then A is better than B, other things being equal.

> *Quantity*: For any pair of positive welfare levels $\boldsymbol{A}$ and $\boldsymbol{B}$, such that $\boldsymbol{B}$ is slightly lower than $\boldsymbol{A}$, and for any number of lives n, there is a greater number of lives m, such that a population of m people at level $\boldsymbol{B}$ is at least as good as a population of n people at level $\boldsymbol{A}$, other things being equal.

> *Quality*: There is a perfectly equal population with very high welfare which is at least as good as any population with very low positive welfare, other things being equal.

Egalitarian Dominance is as uncontroversial as a principle of population ethics gets. Quality is a condition that rules out theories that imply the Repugnant Conclusion. Quantity may have less immediate intuitive appeal than the first two, but it should appeal to those who find some truth in the saying "the more good, the better".

Now, it turns out that the above seemingly week conditions are not mutually compatible, as illustrated by the following simple and well-known impossibility theorem (for risk-free population axiology):

> ***Simple Impossibility Theorem***: There is no population axiology which satisfies Egalitarian Dominance, Quality, and Quantity.[34]

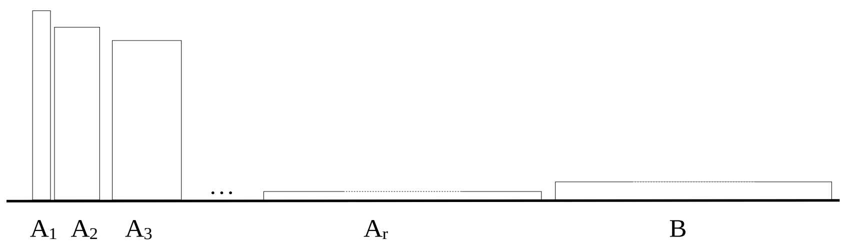

Figure 3.

To see why the above theorem is true, consider the sequence in Figure 3. Assume that A_1 in the diagram above is a population with very high welfare and that B is a population with very low positive welfare (again, the width of the blocks represents the number of lives in the population, the height represents their welfare; dashes indicate that the block in question should really be much wider than shown). According to Quantity, there is a population A_2 with slightly lower welfare than A_1 which is at least as good as A_1; a population A_3 with slightly lower welfare than A_2 which at least as good as A_2; and so forth. Hence, we will finally reach population A_r with very low positive welfare.[35] By transitivity, A_r is at least as good as A_1. But B is better than A_r according to Egalitarian

[34] This theorem was first presented in (Arrhenius 2000b) where exact statements of these conditions also can be found. See also (Arrhenius 2016).

[35] This presupposes that the welfare levels in this sequence of populations satisfy a structural assumption which has been called *Finite Fine-grainedness*: There exists a finite sequence of slight welfare differences between any two welfare levels. The idea behind this weak structural assumption is that one can get from one welfare level to another in a finite number of steps of intuitively slight welfare difference. Examples of such welfare differences could be some minor pain or pleasure or a shortening of life by a minute or two. For a precise definition of "slight welfare difference" and a discussion of possible theories of welfare that violate Finite Fine-grainedness, see (Arrhenius and Rabinowicz 2015; Arrhenius 2016, 2000b, 2011, forthcoming, 2005). For an interesting effort to challenge this structural assumption (in light of the impossibility theorems in population ethics), see (Carlson forthcoming) and (Thomas 2017).

Dominance. Hence, by transitivity, B is better than A_1. However, by Quality, A_1 is at least as good as B. So, we get a contradiction: B is better than A_1 which is at least as good as B. Hence, the Simple Impossibility Theorem: there is no population axiology which satisfies Egalitarian Dominance, Quality, and Quantity.

It is rather easy to see that the Simple Impossibility Theorem can be extended to the following impossibility result for population prospect axiologies, even if we assume that all prospects in the theories domain are non-trivial (in the sense described above):

> ***Simple Prospect Impossibility Theorem***: There is no population prospect axiology which satisfies Egalitarian Dominance, Quality, Quantity, and Population State Dominance.

	s_1	s_2
P_1	C	A_1
P_2	C	A_2
…	…	…
P_r	C	A_r
P_{r+1}	C	B

Table 2.

P_1 to P_{r+1} are different population prospects, whereas s_1 and s_2 are different states of the world that determine which population results from the choice of each prospect. We assume that both the prospects in Table 2 above and those in Table 3 below (Q_1 to Q_{r+1}) are non-trivial, meaning that the probabilities of states s_1 and s_2 are strictly between zero and one.

Now, A_1, A_2, etc., and B in Table 2 are the populations depicted in Figure 3. So, for each A_i, A_{i+1} is at least as good as A_i according to Quantity, B is better than A_r according to Egalitarian Dominance, and thus B is better than A_1 by transitivity. However, A_1 is at least as good as B according to Quality. But that means that according to Population State Dominance, P_{r+1} is better than P_1 and P_1 is at least as good as P_{r+1}, irrespective of what population C represents. So, again we reach a contradiction. Hence, the Simple Prospect Impossibility Theorem: there is no population prospect axiology which satisfies Egalitarian Dominance, Quality, Quantity, and Population State Dominance.

It might be worth reminding the reader of a point that we previously made, namely, that prospect impossibility results like those proven above cannot be achieved by only using risk-free adequacy conditions, such as Egalitarian Dominance, Quantity, and Quality. For while these risk-free conditions tell us how to rank *outcomes*, they tell us nothing about how to rank *prospects* that can result in different outcomes. For instance, while Egalitarian Dominance, Quality, and Quantity together entail that B is better than A_1 and that A_1 is at least as good as B, that does not suffice to get an impossibility result for prospects, since this non-transitive outcome-ranking by itself does not preclude that there exists an adequate (in particular, transitive) population *prospect* ranking.

More generally, to derive a prospect impossibility result, one either has to use a condition like Population State Dominance, to connect the ranking of outcomes to the ranking of prospects; or one has to reformulate some or all of the standard conditions as prospect conditions (a possibility we explore blow). For instance, while Egalitarian Dominance, Quantity, and Quality are mutually inconsistent as conditions on how to rank outcomes, these three conditions are by themselves mutually *consistent* as conditions on how to rank prospects. After all, conditions that say *nothing* about how to rank prospects are obviously mutually consistent as conditions on how to rank prospects.

Another approach, briefly mentioned above, would be to reformulate all of Egalitarian Dominance, Quality, and Quantity as prospect conditions by building Population State Dominance into the original conditions. Consider for instance the following conditions:

> *Egalitarian Prospect Dominance*: For any population C and any probability p, if A is a perfectly equal population of the same size as population B, and every person in A has higher welfare than every person in B, then a prospect that with probability p results in A and with probability 1-p results in C is better than a prospect that with probability p results in B and with probability 1-p results in C, other things being equal.

> *Quantity Prospect Dominance*: For any population C, any probability p, and any pair of positive welfare levels $\boldsymbol{A}$ and $\boldsymbol{B}$, such that $\boldsymbol{B}$ is slightly lower than $\boldsymbol{A}$, and for any number of lives n, there is a greater number of lives m, such that a prospect that with probability p results in a population of m people at level $\boldsymbol{B}$ and with probability 1-p results in C, is at least as good as a prospect that with probability p results in a population of n people at level $\boldsymbol{A}$ and with probability 1-p results in C, other things being equal.

> *Quality Prospect Dominance*: There is a perfectly equal population A with very high welfare such that for any population C, for any probability p, and for any population B with

very low positive welfare, a prospect that with probability p results in A and with probability 1-p results in C is at least as good as a prospect that with probability p results in B and with probability 1-p results in C.

Now, it should be straightforward to see that even if the domain is restricted to non-trivial prospects, that is, prospects that confer a probability strictly between zero and one on different populations, we can prove:

Pure Prospect Impossibility Theorem: There is no population prospect axiology which satisfies Egalitarian Prospect Dominance, Quality Prospect Dominance, and Quantity Prospect Dominance.

The prospects in Table 2, for instance, establish the Pure Prospect Impossibility Theorem. For notice that by transitivity, Egalitarian Prospect Dominance, Quality Prospect Dominance, and Quantity Prospect Dominance result in an inconsistent ranking of the prospects in Table 2, in exactly the same way as the original Egalitarian Dominance, Quality, Quantity and Population State Dominance resulted in an inconsistent ordering of the prospects in Table 2.

It might also be worth pointing out that the assumption that the different prospects result in the *same* population if one of the two states obtains is not essential to the above results. This holds for both the Simple Prospect Impossibility Result and the Pure Prospect Impossibility Result, but to keep things simple, we shall focus on the latter. Consider the prospects in Table 3, where we have extended the sequence in Figure 3 beyond A_r to A_{r+1} and A_{r+2}, and assume that the new population A_{r+1} relates to A_r as A_r relates to A_{r-1}, etc.

	s_1	s_2
Q_1	A_1	A_2
Q_2	A_2	A_3
…	…	…
Q_r	A_r	A_{r+1}
Q_{r+1}	B	A_{r+2}

Table 3.

By Quantity, A_2 is at least as good as A_1 and A_3 is at least as good as A_2. More generally, for each A_i, A_{i+1} is at least as good as A_i. Hence, by Population State Dominance, Q_2 is at least as good as Q_1, and, continuing this sequence, Q_r is at least as good as Q_{r-1}. However, B is better than A_r, so by Population State Dominance, Q_{r+1} is better than Q_r. So, by transitivity, Q_{r+1} is better than Q_1. But by Quality, A_1 is at least as good as B and (we can construct the example such that) A_2 is at least as good as A_{r+2}. Therefore, by Population State Dominance, Q_1 is at least as good as Q_{r+1}. So, again we reach a contradiction: Q_{r+1} is better than Q_1 which is at least as good as Q_{r+1}.[36]

Finally, it might be worth noting that impossibility results can be generated by using only "pure" prospect adequacy conditions; that is, adequacy conditions that involve only (non-trivial) population prospects, but that nevertheless mirror the traditional adequacy conditions involving (risk-free) populations.[37] Consider the following two conditions:

> *Probability/Quantity trade-off*: For any population A consisting of lives with very high welfare, and any population B consisting of lives with very low but positive welfare, and any prospect P_1 that has probability $p>0$ of resulting A and probability 1-p of resulting in B, there is some probability $q<p$, and some population C consisting of lives at the same welfare level as the B-people, such that a prospect P_2 that has probability q of resulting in population A and probability 1-q of resulting in the population consisting in the union of B with C, is no worse than prospect P_1, other things being equal.

> *Probabilistic Average/Total trade-off*: There is a perfectly equal population A with very high welfare, such that for any two populations B and C consisting of lives with the same very low but positive welfare, there are some probabilities p, q, where $q<p$, such that prospect P_1 that has probability p of resulting in population A and probability 1-p of resulting in B is better than prospect P_2 that has probability q of resulting in population A and probability 1-q of resulting in the union of B with C, other things being equal.

---

[36] We should also note that an analogous, and equally simple, prospect impossibility result can be derived by combining the above three traditional adequacy conditions (i.e., Quantity, Quality, and Egalitarian Dominance) with *Stochastic Dominance*, which can be stated as follows:

> *Stochastic Dominance*: If for any outcome x prospect P_1 has at least as high a chance as prospect P_2 of resulting in an outcome that is at least as good as x, then P_1 is at least as good as P_2; and if, in addition, for some outcome y prospect P_1 has a higher chance than prospect P_2 of resulting in an outcome that is better than y, then P_1 is better than P_2.

In a framework like the one we are considering (in particular, one that contains states of the world), Stochastic Dominance entails State-Wise Dominance. Hence, since we cannot consistently add Population-State Dominance to the three traditional adequacy conditions discussed above, we also cannot consistently add Stochastic Dominance to these three adequacy conditions.

[37] We thank Tim Campbell for suggesting a result like this.

The thought behind Probability/Quantity trade-off is that we can reduce the probability that a prospect results in the population A, where people have very high welfare, without making the resulting prospect worse, as long as we sufficiently increase the number of people in the alternative population that the prospect results in if it does not result in A.[38] In contrast, the thought behind Probabilistic Average/Total trade-off is that there is a limit to how much average welfare we are willing to risk foregoing for the sake of a chance at higher total welfare. As can easily be verified, Prospect Average Utilitarianism violates Probability/Quantity trade-off whereas Prospect Total Utilitarianism violates Probabilistic Average/Total trade-off.

To see how Probability/Quantity trade-off and Probabilistic Average/Total trade-off jointly lead to inconsistency, given transitivity, suppose we start with a probability p that is close to 1. By Probability/Quantity trade-off, there is some quantity δ that we could subtract from p such that the following holds. Suppose that prospect P_0 has probability p of resulting in a population A consisting of lives with very high welfare and probability 1-p of resulting in some population B consisting of lives with very low but positive welfare. Moreover, suppose that prospect P_1 has probability p-δ of resulting in population A and probability 1-(p-δ) of resulting in the union of B with some population C consisting of lives with the same low welfare as the B-lives. Then P_1 is no worse than P_0. But continuing this reasoning, we can generate a sequence of prospects, by incrementally increasing the size of the quantity that is subtracted from p, until we reach a quantity ζ for which the following holds. Suppose that P_n has probability p-ζ of resulting in A and probability 1-(ζ) of resulting in some population D, which consists of lives with the same low but positive welfare as the B-lives but which is much bigger than B (and indeed contains both B and C). Then P_n is no worse than the immediately preceding prospect in the sequence. Hence, by transitivity, P_n is no worse than P_0.[39] Since we get this result irrespective of the welfare (and size) of the A population, and irrespective of p and ζ, we have a violation of Probabilistic Average/Total trade-off, according to which P_0 is better than P_n (for some A, p, ζ, and D).

---

[38] One could weaken the principle by excluding the ends of the probability spectrum, i.e., by adding that q is not equal to 0 and p is not equal to 1. Weakening the principle in this way would still generate the inconsistency with Probabilistic Average/Total trade-off (assuming transitivity). But note that Prospect Total Utilitarianism entails the stronger version, where the possible values of p and q are not restricted.

[39] Those who worry about incomparability (or incommensurability), might question this argument, since the relation of being no worse than is not transitive when applied to incomparable options. To get around this worry, we could strengthen Probability/Quantity trade-off such that it says that P2 is *better than* P1. The obvious cost of strengthening the principle is that some might only accept the weaker version of it. However, the general motivation for the weaker principle arguably supports the stronger principle too. In any case, incomparability (and incommensurability) raises various interesting questions in population ethics that are beyond the scope of this paper.

V. Concluding Remarks

In conclusion, we have found that there can be no adequate population prospect axiology, since no such theory can satisfy conditions that, we contend, any adequate population prospect axiology should satisfy. And although we have focused only on five such adequacy conditions, and only one of the many impossibility theorems that have been proven for risk-free population axiology, it should be evident that the other theorems can be extended to population prospect axiology in an analogous way to how we extended the Simple Impossibility Theorem, i.e., by exploiting Population State Dominance.

On one hand, the result of this paper should perhaps not come as a surprise. Given that we would want the traditional adequacy conditions from risk-free population axiology to hold when ranking the (risk-free) outcomes (i.e., populations) in population prospects, and since we in addition want to impose further adequacy conditions that have specifically to do with the introduction of uncertainty, we should perhaps not have expected the search for an adequate population prospect axiology to be easier than the search for an adequate population axiology. On the other hand, one might have hoped that the impossibility theorems from risk-free population ethics could be avoided by removing some of their idealizing assumptions, for instance, the assumption that we know for sure which population an alternative will result in. But, as we now have seen, removing the idealizing assumption of certainty will not help us avoid these impossibility results.[40]

Bibliography

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.

Arrhenius, Gustaf. forthcoming. 'Population Ethics: The Challenge of Future Generations'. Oxford University Press.

———. 2000a. 'An Impossibility Theorem for Welfarist Axiologies'. *Economics and Philosophy* 16 (02): 247–266.

———. 2000b. *Future Generations: A Challenge for Moral Theory*. Uppsala: University Printers. http://www.diva-portal.org/smash/record.jsf?pid=diva2:170236.

———. 2005. 'Superiority in Value'. *Philosophical Studies* 123 (1/2): 97–114. https://doi.org/10.1007/s11098-004-5223-0.

———. 2011. 'The Impossibility of a Satisfactory Population Ethics'. In *Descriptive and Normative Approaches to Human Behavior, Advanced Series on Mathematical Psychology*, edited by Hans Colonius and Ehtibar N. Dzhafarov, 1–26. World Scientific Publishing Company.

[40] We would like to thank Matt Adler, Geir Asheim, Krister Bykvist, Tim Campbell, Robert Goodin, Julia Mosquera, Jake Nebel, Wlodek Rabinowicz, Daniel Ramöller, Dean Spears, Alex Voorhoeve, and Stéphane Zuber for very helpful discussions. Thanks also to the audiences at the *New Scholarship on Population Ethics*, Duke University, April 24, 2017; *Sustainability and Population Ethics*, Iméra Aix Marseille Université, 3 July, 2017; *Boulder Formal Value Theory Workshop*, March 19; and the workshop *Population Ethics*, Institute for Futures Study, May 4, 2017 for useful questions and comments. Financial support from the Swedish Research Council and Riksbankens Jubileumsfond is gratefully acknowledged.

———. 2016. 'Population Ethics and Different-Number-Based Imprecision'. *Theoria* 82 (2): 166–81. https://doi.org/10.1111/theo.12094.

Arrhenius, Gustaf, and Wlodek Rabinowicz. 2015. 'Value Superiority'. In *The Oxford Handbook of Value Theory*, edited by Iwao Hirose and Jonas Olson, 225–48. New York NY: Oxford University Press.

Asheim, Geir B., and Stéphane Zuber. 2016. 'Evaluating Intergenerational Risks'. *Journal of Mathematical Economics* 65 (August): 104–17. https://doi.org/10.1016/j.jmateco.2016.05.005.

Bales, R. Eugene. 1971. 'Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?' *American Philosophical Quarterly* 8 (3): 257–265.

Bell, David E. 1982. 'Regret in Decision Making under Uncertainty'. *Operations Resarch* 30(5): 961-981.

Blackorby, Charles, Walter Bossert, and David J. Donaldson. 2005. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. New York: Cambridge University Press.

Brink, David O. 1986. 'Utilitarian Morality and the Personal Point of View'. *The Journal of Philosophy* 83 (8): 417–38.

Budolfson, Mark, and Dean Spears. Mimeo. 'All Leading Population Axiologies Imply the Repugnant conclusion'.

Broome, John. 1991. *Weighting Goods*. Oxford: Oxford University Press.

———. 1992. *Counting the Cost of Global Warming*. Cambridge: The White Horse Press.

———. 2004. *Weighing Lives*. Oxford ; New York: Oxford University Press.

———. 2010. 'The Most Important Thing about Climate Change'. *Why Ethics Matters*, 101.

———. 2015. 'Climate Change: Life and Death'. In *Climate Change and Justice*, edited by Jeremy Moss, 184–200. Cambridge University Press.

Carlson, Erik. forthcoming. 'On Some Impossibility Theorems in Population Ethics'. Mimeo. Department of Philosophy, Uppsala University.

Danielsson, Sven. 1974. 'Two Papers on Rationality and Group Preferences'. *Filosofiska Studier*, Uppsala: Filosofiska föreningen och filosofiska institutionen, 21.

Diamond, Peter A. 1967. 'Cardinal Welfare, Individualist Ethics, and Interpersonal Comparison of Utility: Comment'. *Journal of Political Economy* 75(5): 765-766.

Diecidue, Enrico, and Jeeva Somasundaram. 2017. 'Regret Theory: A new Foundation'. *Journal of Economic Theory* 172(2): 88-119

Kahneman, Daniel, and Amos Tversky. 1979. 'Prospect Theory: An Analysis of Decision under Risk'. *Econometrica* 47(2): 263-291.

Nebel, Jake. 2017. 'Priority, not Equality, for Possible People'. *Ethics*, 127(4) 896-911.

Nebel, Jake. 2019. 'An Intrapersonal Addition Paradox'. *Ethics,* 129(2): 309-343

Loomes, Graham, and Robert Sugden. 1982. 'Regret Theory: An Alternative Theory of Rational Choice under Uncertainty'. *The Economic Journal* 92(368): 805-824.

Parfit, Derek. 1984. *Reasons and Persons*. 1991st ed. Oxford: Clarendon.

Quiggin, John. 1982. 'A Theory of Anticipated Utility'. *Journal of Economic Behavior and Organization* 3(4): 323-343.

Roberts, Melinda A. 2007. 'The Non-Identity Fallacy: Harm, Probability and Another Look at Parfit's Depletion Example'. *Utilitas* 19 (03): 267–311.

Stefánsson, H. Orri. 2015. 'Fair Chance and Modal Consequentialism'. *Economics and Philosophy* 31(3): 371-395.

Stefánsson, H. Orri, and Richard Bradley. 2015. 'How Valuable are Chances?' *Philosophy of Science* 82(4): 602-605.

Stefánsson, H. Orri, and Richard Bradley. 2019. 'What is Risk Aversion?' *British Journal for the Philosophy of Science* 70(1): 77-102.

Thomas, Teruji. 2017. 'Some Possibilities in Population Axiology'. *Mind*, December. https://doi.org/10.1093/mind/fzx047.

Tversky, Amos and Daniel Kahneman. 1992. 'Advances in Prospect Theory: Cumulative Representation of Uncertainty'. *Journal of Risk and Uncertainty* 5(4): 297-323.

Voorhoeve, Alex, and Marc Fleurbaey. 2016. 'Priority or Equality for Possible People?' *Ethics*, 126 (4): 929–954.