Forthcoming in *The Journal of Philosophy*

Path Independence and a Persistent Paradox of Population Ethics

Rush T. Stewart *

Abstract

In the face of an impossibility result, some assumption must be relaxed. The Mere Addition Paradox is an impossibility result in population ethics. Here, I explore substantially weakening the decision-theoretic assumptions involved. The central finding is that the Mere Addition Paradox persists even in the general framework of choice functions when we assume Path Independence as a minimal decision-theoretic constraint. Choice functions can be thought of either as generalizing the standard axiological assumption of a binary "betterness" relation, or as providing a general framework for a normative (rather than axiological) theory of population ethics. Path Independence, a weaker assumption than typically (implicitly) made in population ethics, expresses the idea that, in making a choice from a set of alternatives, the order in which options are assessed or considered is ethically arbitrary and should not affect the final choice. Since the result establishes a conflict between the relevant ethical principles and even very weak decision-theoretic principles, we have more reason to doubt the ethical principles.

Keywords. Axiology; impossibility theorem; mere addition paradox; path independence; population ethics; rational choice

1 Introduction

In both policy-making and private life, we face choices that affect the (expected) makeup and welfare of future generations. What are optimal abatement strategies in the context of climate change? Should a professional couple in their forties try to have a child? Should one take up hang gliding as a hobby at a young age? A central task in population ethics is the articulation of a theory for assessing the various populations and welfare distributions that such choices might influence. An *axiology* is a type of binary "betterness" relation that compares different populations with different welfare profiles, and is typically taken as focal in developments of population ethics (Greaves, 2017, p. 1; Thomas, 2018, p. 811). The axiological data, even if not fully determinative, are taken to be important considerations for decisions that impact future generations.

*Thanks to Jean Baccelli, Stephan Hartmann, Christian List, Jorge Morales, Michael Nielsen, Ignacio Ojea Quintana, Reuben Stern, Mark Swails, and audiences at King's College London and LMU Munich for helpful comments and discussions. Thanks to an anonymous referee for a very helpful report, and for a number of suggestions resulting in the final section of the paper. I'm grateful to the Center for Advanced Studies at LMU Munich and Longview Philanthropy for providing me the freedom and time to pursue this project. The diagrams were made at www.boxethics.org.

A number of "paradoxes" arise in population axiology. Certain sets of apparently desirable properties of an axiology are impossible to jointly satisfy. A prominent instance of this is the *Mere Addition Paradox* (Parfit, 1984). Two tempting axiological judgments—and thus the initially compelling axiological principles that imply them—entail the so-called *Repugnant Conclusion*: for any population consisting of people who are thriving, there is another (much larger) population consisting of people with lives barely worth living that is better. As the name indicates, this is a conclusion many hope to avoid.

Confronted with limitative results in population axiology, one response is to surrender some of the substantive, ethical desiderata. For instance, some are able to reconcile themselves with the Repugnant Conclusion (e.g., Tännsjö, 2002; Huemer, 2008). I will not consider this strategy here, though one response to the central result below may be to increase interest in it. Another response is to relax some of the more "structural" or decision-theoretic assumptions (see Greaves, 2017, p. 12). Temkin, for example, advocates relaxing transitivity of the betterness relation (1987). Thomas, who entertains forms of incompleteness, considers sacrificing certain other things "less extreme" (2018, p. 822) and "less painful" than giving up transitivity (2018, p. 829), agreeing with Broome that transitivity is "an analytic feature" of comparatives (2018, p. 810). In his Dewey Lectures, Sen challenges the imposition of certain structural assumptions at the outset of ethical theorizing: "Yielding complete orders cannot be an a priori requirement of the legitimacy of a moral principle, or even of an entire moral approach with its full structure of principles" (1985, p. 180). This is of a piece with a point Sen makes repeatedly in his work: much of the work we need social choice theory to do can be done without completeness and certain impossibility results can be avoided by dropping it. Extending Sen's point, we might challenge even the presumption of any sort of binary relation as overly strong apriorism. Perhaps moral principles, or even entire moral approaches with their full structures of principles, do not reduce to binary comparisons; perhaps even binary relations are too restrictive a basis for ethical choice in general.

Rather than begin with the assumption of a binary relation, I take general abstract choice theory as a setting for investigating some issues in population ethics. This is a move familiar from social choice theory for exploring potential escape routes from impossibility theorems. Choice functions select a subset of a set of available alternatives. The selected alternatives can be interpreted as the "good" or "permissible" alternatives. Since not all choice functions can be thought of as maximizing a binary relation let alone a binary relation of a particular sort, choice functions are a substantial generalization of assuming an axiology. According to Greaves's survey of the field, some "suggest that one might escape from the impossibility theorems [...] by appealing to vagueness or incompleteness of the betterness relation" (Greaves, 2017, p.12). In the framework presented here, various forms of vagueness or incompleteness are allowed, and I explain a natural one in Section 2 and a generalization of it in Section 5. So abstract choice theory would seem to be a framework for investigating possibilities and limitations in population ethics in a very general way. Section 2 introduces and motivates the framework of abstract choice theory. Some strenuously resist the idea that a theory of goodness or value admits generalization beyond binary relations (e.g., Broome, 2009b). While I indicate some ways to think of choice functions as capable of delivering such a theory below, other interpretations are available. In particular, choice functions can also be interpreted as articulating a normative theory of population ethics, as selecting options that are permissible or consistent with moral obligations.

Absent additional constraints on choice, choice functions are *too* flexible. The bare fact that some set of intuitions can be accommodated in a formal framework does not by itself vindicate those intuitions. The main constraint that I consider is *Path Independence*.[1] Compared to standard assumptions in population ethics, this is a rather weak assumption. Path independent choice functions are not generally determined by binary choice data. There may exist no binary relation that "rationalizes" a path independent choice function. One way to think of Path Independence in the context of ethical decisions is as encoding the idea that, since the order of consideration or assessment of alternatives in a set is arbitrary from an ethical point of view, the goodness or permissibility of alternatives does not depend on it. Greaves continues her survey of responses to impossibility theorems in population ethics: "Still others argue that the normative force of the impossibility theorems depends on an assumption of choice-set-independence, and then go on argue that the choice-set-dependence exhibited by [certain theories] may be positively a virtue of those theories, so that we can hope to find an adequate theory in this space" (2017, p. 12). As we will see, *Path* Independence is consistent with forms of *context* dependence including violations of what Greaves calls choice-set-independence. Consider, for example, if our professional couple in their forties faces a choice like one that Parfit discusses (Parfit, 1996, p. 311). In a choice between having no child and having a child with a mild disability, at least under certain assumptions, both choices are permissible. If there is an additional option to have the same child without the disability, it might be plausible to think that choosing to have a disabled child would not be permissible in this three-way choice, but it remains permissible to choose not to have a child. That the permissibility of choosing to have a child with a mild disability in the presence of the option not to have a child depends on what other options are available is a form of context dependence (and a violation of choice-set-independence) that is verboten according to many theories. Yet, it is perfectly consistent with the assumption of Path Independence as I explain in more detail in Section 3. Allowing for this sort of context dependence might count as a motivation for Path Independence.

However, in Section 4, I show that the Mere Addition Paradox persists if we assume even the weak property of Path Independence. Put another way, avoiding the Mere Addition Paradox while retaining the substantive ethical judgments ingredient to it entails path dependence, that the order in which alternatives are assessed makes a difference to which options are good or permissible. From results in rational choice theory, this is equivalent to violating at least one of two weak and compelling choice consistency properties, as I explain in Section 2. The fact that the Mere Addition Paradox persists even here places limits on the shape that the optimism of those willing to reject choice-set-independence can take. There are at least two ways of interpreting the result, depending on one's other theoretical commitments. If, like Arrhenius (2004) for example, one is open to the prospect of generalizing a theory of value or goodness beyond binary relations, then I state a general limitative result for a theory of value or goodness in population ethics.[2] If one agrees with Broome (e.g., 2009b)

[1] Path Independence enjoys considerable theoretical backing, finding application and motivation in social choice theory (e.g., Arrow, 1963; Plott, 1973), matching theory (e.g., Chambers and Yenmez, 2017), and decision theory for imprecise probabilities (e.g., Levi, 1980) as well as in more abstract areas (e.g., Johnson and Dean, 2001; Koshevoy, 1999; Monjardet and Raderanirina, 2004; Danilov and Koshevoy, 2005; Bossert et al., 2009).

[2] Arrhenius writes, for example, "there are ways of understanding value-concepts such as 'better than' in terms of normative ones such as using "A is better than C" as synonymous with 'A ought to be chosen in a

and others that a theory of goodness or value requires a unique axiology, the result I state can be interpreted as a new impossibility theorem for normative theories of population ethics. As Thomas suggests, "It is at least arguable that the same problems arise if we skip axiology and deal directly with the question of which populations one ought to bring about" (2018, p. 808, fn. 3). On this second reading, I provide some confirmation that worries in population axiology extend to normative theories even under weak assumptions. In Section 5, I consider a potential objection to Path Independence.

2 Choice and Permissibility

Since the points I want to make here do not require a lot of technical machinery, I will work in a minimal framework (cf. Thomas, 2018). I assume that lives vary in how satisfying they are to the individuals living them. At least some lives go better than others.[3] If two lives go equally well for the people living them, they have the same welfare level. Standardly in population ethics, a particular welfare level (or set of levels) is singled out as neutral. Welfare levels above it correspond to lives worth living; welfare levels below correspond to lives not worth living. Moreover, among welfare levels above neutrality, some correspond to "drab" lives, lives that are "barely worth living." A disjoint set of the welfare levels above neutrality correspond to "blissful" lives, lives that go very well.[4] A profile A is a finite, unordered list of welfare levels. Any population of lives determines an associated profile. I will use letters from the beginning of the alphabet for profiles. If A and B are profiles, let $A \cup B$ be the profile that concatenates the profiles A and B.

The present study takes choice functions rather than binary relations as primitive. Let $\mathbb{A}$ be the (nonempty) set of all relevant profiles. A *menu* is a finite, nonempty subset of $\mathbb{A}$. I will use letters towards the end of the alphabet for menus. A *choice function* C associates any menu S with a nonempty subset $C(S)$ of S called the *choice set*. In rational choice theory, the choice set is interpreted as the set of *acceptable* or *admissible* options in the menu.[5] A choice function is *rationalizable* by a binary relation if there exists a relation R

situation where A and C are the only alternatives,' 'A is more choice-worthy than C,' and the like. From this understanding of 'better than,' no axiological Mere Addition Paradox follows since we have explained 'better than' in terms of normative concepts and these claims are restricted to pairwise comparisons and there is no plausible analogue to the transitivity of 'better than' for these normative concepts" (2004, pp. 207-208). The central result that I present is a choice-theoretic Mere Addition Paradox.

[3] We can assume a preorder—a reflexive and transitive binary relation—on the set of welfare levels (cf. Thomas, 2018, p. 811), though, in principle, there may be reasons to generalize this assumption, too.

[4] All of this assumes some highly non-trivial things about the measurability and comparability of welfare, but I will follow the literature in population ethics and regard such assumptions as "legitimate idealizations" for the purposes of this essay (Carlson, 1998, p. 284). Some are not willing to grant as much. For example, Narens and Skyrms write, "Parfit's argument [regarding the Repugnant Conclusion] is meaningless for interval-scale utility. For this reason, modern decision theorists do not think much of Parfit's argument" (2020, pp. 153-54).

[5] I am assuming that choice functions are defined for *any* menu. While this assumption is quite standard, important work has considered generalizing rational choice to arbitrary domains (e.g., Suzumura, 1983; Chambers and Echenique, 2016). By assuming universal domain (over finite menus), I am assuming that a theory of population ethics partitions any finite set of options into those that are permissible—the choice set—and those that are not. I am also assuming that choice functions return a *nonempty* subset for each menu. The most natural interpretation of a menu is as the set of all feasible options. On this interpretation, to relax the stipulation that choice sets are nonempty is to allow that there are cases in which *no* feasible option is permis-

on $\mathbb{A}$ such that, for any menu S, $C(S)$ is the set of R-maximal options in S.[6] A standard assumption about the relation being maximized is that it weakly orders the alternatives, that is, it is a complete and transitive binary relation (Von Neumann and Morgenstern, 1944; Arrow, 1963; Savage, 1954). In this case, the choice function generated by maximizing such a relation would be *weak order* rationalizable. Both types of choice function can be characterized in terms of simple "consistency" constraints on the choice function.[7] But not all choice functions are weak order rationalizable or even rationalizable. Various other properties have been extensively explored, and one plays a central role in this study. As I discuss below, different interpretations of $C(S)$ are available. As a generalization of the notion of an axiology, $C(S)$ might be interpreted as the good options on the menu or the value-admissible options (or, given Path Independence, as the set of best options according to the *set* of legitimate axiologies, as I explain below). As a primitive for a normative or deontic theory, $C(S)$ can be interpreted as selecting the permissible options on menu S. The relevant notion of permissibility has to be restricted to purely population ethics concerns insofar as there are other determinants of moral obligation and permissibility besides the considerations deriving from population ethics. I will use *permissible* generally to describe the options in $C(S)$ unless I am specifically discussing different interpretations of choice sets.

A generalization of weak order rationalizability of considerable, independent theoretical concern is Path Independence (Arrow, 1963; Plott, 1973; Chambers and Yenmez, 2017). Path Independence allows larger choice problems to be decomposed into smaller ones. The choice from the larger menu will not depend on the choice path through the sub-menus.

> **Path Independence.** The permissible options in the menu $S \cup T$ are the same as the permissible options from the menu consisting of the permissible options of S and the permissible options of T.[8]

Interpreting choice as selecting the permissible profiles from the feasible ones, Path Independence asserts that the permissible options do not depend on the order of assessment or consideration of the alternatives. If the order of assessment or consideration of options is ethically arbitrary for a theory of population ethics, Path Independence seems like a desirable feature. Failure of Path Independence would also leave a decision maker open to forms of manipulation when the order of presentation can be influenced.

sible, which would seem to run into violations of ought-implies-can for normative theories of population ethics and, when the choice function is rationalizable by a binary relation or a set of them, violations of acyclicity for axiological theories. To treat concern about nonempty choice sets explicitly, some choice theorists stipulate that a "status quo" or "abstain" option is present in every menu (e.g., Fishburn, 1973)—though, one might complain, if that option is feasible, it should already be included. But, as with general choice domains, there is work allowing for empty choice sets (e.g., Aizerman, 1985). One could investigate the extent to which certain stock rational choice assumptions can be relaxed without altering the analysis here, including the central observation in particular.

[6] That is, there exists an $R \subseteq \mathbb{A} \times \mathbb{A}$ such that, for all menus S, $C(S) = \{A \in S : \neg\exists B \in S \text{ such that } BPA\}$, where P is the asymmetric part of R, i.e., APB if and only if ARB and not BRA. If R is complete, then the sets of R-maximal and R-optimal elements coincide: for all menus S, $\{A \in S : \neg\exists B \in S \text{ such that } BPA\} = \{A \in S : \forall B \in S\ ARB\}$.

[7] Rationalizability simpliciter is characterized by the conjunction of Property α, stated in the body of the essay, and Property γ: for all nonempty $S, T \subseteq \mathbb{A}$, $C(S) \cap C(T) \subseteq C(S \cup T)$. Rationalizability by a weak order is characterized by the conjunction of Property α and Property β: for all nonempty $S, T \subseteq \mathbb{A}$, if $S \subseteq T$, $A, B \in C(S)$ and $A \in C(T)$, then $B \in C(T)$.

[8] For all nonempty $S, T \subseteq \mathbb{A}$, $C(S \cup T) = C(C(S) \cup C(T))$. When applied to menus or sets of profiles, union has its usual set-theoretic meaning.

Here are two alternative ways to think about Path Independence. First, Path Independence is equivalent to the conjunction of two standard and compelling axioms (Aizerman and Malishevski, 1981; Moulin, 1985, Lemma 6). The first is Property α, "the mother of all choice consistency conditions" (Nehring, 1997, p. 407).

> **Property** α. If A is permissible in a menu T, then A is permissible in any subset of T that contains it.[9]

In other words, a permissible option cannot be made impermissible by removing other alternatives from the menu. Property α plays a very central role in abstract choice theory.[10] It is a necessary condition for weak order rationalizability, rationalizability simpliciter, and pseudo-rationalizability (defined below) among other important concepts. The second condition is Aizerman's Axiom (a slight weakening of a condition called "independence of rejecting the outcast variants").

> **Aizerman's Axiom.** If all the permissible alternatives in a menu T are included in a subset S of T, then the only permissible alternatives in S are ones that are permissible in T.[11]

In other words, if all of the permissible options in a big menu T are available in a smaller menu S, then all the permissible options in S are ones that are permissible in the larger menu T. Where Property α says that permissible options remain permissible in smaller menus, Aizerman's Axiom says that options that are not permissible in a menu do not become permissible by removing some non-permissible options from the menu. Path Independence is equivalent to the conjunction of Property α and Aizerman's Axiom, so the attractiveness of Path Independence can be subjected to more refined intuitive evaluation by considering the two axioms separately. Any violation of Path Independence entails a violation of at least one of these two properties.

While my focus is mostly on the choice-theoretic constraints and framing, a second way of thinking about Path Independence in the context of population ethics involving *indeterminacy* is available. Aizerman and Malishevski establish that path independent choice functions have a very nice decomposition (Aizerman and Malishevski, 1981, Theorem 3; Moulin, 1985, Theorem 5).[12] If C is path independent, then there exists a *set* of weak orderings $\{\succsim_i\}_{i \in I}$ such that, for every menu S, $C(S) = \bigcup_{i \in I} M(S, \succsim_i)$, where $M(S, \succsim_i)$ is the set of $\succsim_i$-maximal elements of S.[13] When C has such a decomposition, C is called *pseudo-rationalizable*. So, Path Independence is equivalent to pseudo-rationalizability. In this case, we could say that

[9]For all nonempty $S \subseteq T \subseteq \mathbb{A}$, $C(T) \cap S \subseteq C(S)$.

[10]While Property α enjoys very wide support, Sen, for example, has offered some potential counterexamples (e.g., Sen, 1994, pp. 250-251). However, various attractive and interesting ways of accommodating them and recovering the standard theory as a special case have been developed (e.g., Bossert and Suzumura, 2009; Stewart, 2016; Helzner, MS). In any event, Property α is common ground with any approach to population ethics given in terms of an axiology on $\mathbb{A}$.

[11]For all nonempty $S, T \subseteq \mathbb{A}$, $C(T) \subseteq S \subseteq T$ implies $C(S) \subseteq C(T)$.

[12]These proofs are for finite $\mathbb{A}$, but see (Pedersen, 2009). But inspection of Moulin's Theorem 5, for example, reveals that the equivalence of the Path Independence, on the one hand, and Property α and Aizerman's Axiom, on the other, does not depend on the finiteness assumption.

[13]For any path independent choice function, there also exists a (typically larger) set of linear orders and a set of preorders—weak orders are preorders—either of which could be substituted for the weak orders in the statment of pseudo-rationalizability.

there is a set of legitimate axiologies rather than a unique one, and C selects the options that are best according to *some* legitimate axiology. Legitimate axiologies might be thought of as representing distinct but permissible ways of aggregating or trading off various dimensions of value—like *general welfare* and *equality* (cf. Thomas, 2018, p. 822, fn. 27). When there are multiple legitimate axiologies, there may fail to be an "determinately" best option (cf. Broome, 2009a; Rabinowicz, 2009). For example, it could be the case that $A \succ_1 B$ and $B \succ_2 A$. In that case, it is indeterminate whether A or B is better. But if $A \succ_i B$ for all $i \in I$, then it's fair to say that A is determinately or categorically better than B. There are a number of ways to define a categorical axiology from a set of legitimate axiologies. The simplest is to just take the intersection of the legitimate axiologies, which yields a preorder $\succsim = \bigcap_{i \in I} \succsim_i$.[14] Two options might be comparable according to any legitimate axiology, but be incomparable according to the corresponding categorical relation. But, to reiterate, an important point about path independent choice functions is that there may be *no* rationalizing binary relation, including the categorical one.[15]

This form of indeterminacy might be especially interesting to population ethicists for two reasons. First, utilitarianism is a popular view in population axiology. Two prominent variants are Totalism and Averagism. Totalism maintains that A is better than B if and only if total welfare is higher in A than in B, and A and B are equally good if and only if they are equal in total welfare. According to Averagism, A is better than B if and only if average welfare in A is higher than in B, and A and B are equally good if and only if they are equal in average welfare (Greaves, 2017, p. 2). While Totalism and Averagism agree for populations of the same size, they can disagree in the context of variable population sizes. But what if *individual* preference or welfare is indeterminate? Work in economics (e.g., Dubra et al., 2004) and philosophy (e.g., Levi, 1986) has proposed "multi-utility" representations of incomplete preferences. How should we compute totals and averages of utility for such cases? One natural idea might be to compute sets of totals and averages, transferring the indeterminacy to the social level. Second, given both the theoretical difficulties involved in settling on a particular, concrete population axiology and the need to take actions before a satisfactory resolution is achieved, we are often forced to act under *moral* and *axiological* uncertainty (Greaves and Ord, 2017). As I suggested above, Path Independence can be interpreted as selecting the options that are best according to some *set* of legitimate axiologies (without assuming that choice is necessarily governed by a so-called *effective axiology*). A legitimate axiology might be one of the proposals currently under contention like Totalism, or some sort of plausible compromise between different proposals about which axiologists are uncertain like Averagism and Totalism. To take just one example, such a compromise might be a ranking of expected axiological value derived from assigning probabilities to different live axiologies—again tabling thorny issues about intertheoretic comparability—and one reason there may be a set of such compromises could be that probabilities are imprecise. But many

[14] Other proposals include $\succsim = \bigcap_{i \in I} \succ \cup \bigcap_{i \in I} \sim_i$ (e.g., Sen, 2004, p. 672) and $(\succ, \sim, \succsim) = (\bigcap_{i \in I} \succ_i, \bigcap_{i \in I} \sim_i, \bigcap_{i \in I} \succsim_i)$ (e.g., Levi, 2008). Notice that these various relations inherit transitivity from the legitimate axiologies.

[15] Here is a simple example. Let $\mathbb{A} = \{A, B, D\}$. Suppose C is pseudo-rationalized by $\{\succsim_1, \succsim_2\}$, where $A \succ_1 B \succ_1 D$ and $D \succ_2 B \succ_2 A$. The "second place" option B is not optimal relative to any legitimate axiology (which, we can assume, already trade off relevant dimensions of value). The intersection of the legitimate axiologies is empty in this case. But $C(\{A, B, D\}) = \{A, D\}$ even though no option is strictly better than B according to the categorical axiology.

other interpretations of the set of rankings are possible.

My intention here is not solely to promote Path Independence, although it strikes me as pretty compelling and essentially common ground in this context. Even if Path Independence is ultimately found not to be an appropriate standard, I think the general setting of abstract choice theory—with its various axioms, constructions, and results—has potential as a framework for considering issues in population ethics. Here is a summary of some reasons to consider adopting choice functions as primitive in investigations in population ethics. First, choice functions generalize the framework of maximizing binary relations. So, abstract choice theory is a more general and flexible setting to carry out investigations in population ethics. For example, by making weaker, more general assumptions, abstract choice theory offers the prospect of stating stronger, more informative impossibility results. And there are many results about choice functions that may facilitate ethical theorizing. Using choice functions rather than assuming a particular type of ordering out of the gate is a strategy that has been explored in social choice theory (e.g., Sen, 1977).[16] Impossibility results in social choice theory spurred the search for more general frameworks, which in turn resulted in more general and refined analyses of (im)possibilities for social choice.

Second, and relatedly, various forms of indeterminacy can be simply expressed and characterized. This is the case with the pseudo-rationalizability representation of path independent choice. For Path Independence, not only are there various qualitative axioms to scrutinize, there is an interesting construction to contemplate the relevance of for population ethics (though, one might use path independent choice without any appeal to or interest in the associated set of orderings). For example, pseudo-rationalizability suggests one natural interpretation of incompleteness in the categorical betterness relation as conflict between legitimate axiologies. Conflict in values is a possibility pointed out by many value pluralists. Even with a(n incomplete) categorical betterness relation at hand, the connection between it and permissible choices may be more subtle than simply choosing the maximal options with respect to it. Again, path independent choice functions may fail to be rationalizable by any binary relation.

A third motivation for using choice functions is that, even if, like Broome (2009b), one insists on binary betterness relations for studying goodness or value, choice functions are a natural setting to study normative or "deontic" accounts of population ethics, that is, accounts of permissibility or consistency with moral obligation. (For my part, I think a choice-theoretic account can be interpreted as evaluative, but not necessarily in terms of a (single) binary relation in general.[17]) The ultimate point of a population axiology might be to inform such a normative theory. As I explain just below, the strategy of using choices rather than comparisons to frame a normative theory has been suggested a few times in the literature.

A brief overview of some related literature may help to contextualize the contribution of

[16] As Plott (1973, p. 1075) points out, Arrow (1963, p. 120) appeals to Path Independence in response to critics of his "rationality" assumption of a complete and transitive binary social preference relation. But as he shows and is now widely appreciated, Path Independence is implied by but does not imply weak order rationalizable choice.

[17] At least not in terms of a binary relation on $\mathbb{A}$. Non-binary choice functions, including path independent ones in particular, can be represented in terms of binary *hyper*-relations on $2^{\mathbb{A}}$ (Nehring, 1997; Stewart, 2020a). Hyper-relations compare, not just options (as singleton menus), but menus of options also. In fact, hyper-relation transitivity is the central condition in the hyper-relation characterization of Path Independence (Nehring, 1997, Theorem 1.ii, Theorem 6).

this study. The first to show that the Mere Addition Paradox follows from the axiological versions of the particular ethical assumptions used in Section 4 was Ng (1989). Carlson (1998) and Huemer (2008) follow Ng in pointing out that these assumptions are inconsistent. All of these derivations assume a binary relation that is transitive in (at least) its asymmetric factor. In contrast, I do not assume a relation. Boonin-Vail (1996) argues that the judgments ingredient to the Mere Addition Paradox are consistent when phrased in terms of choices from binary choice contexts. The Oughtness Paradox—as Boonin-Vail calls it—for normative theory, unlike the original Goodness Paradox for axiology, "can be solved" (Boonin-Vail, 1996, p. 279). Like Boonin-Vail, I appeal to choice rather than comparatives. I explicitly invoke abstract choice theory as a formal setting and explore one plausible constraint on such permissible choices. But it is readily apparent that not all choices we face are between just two options. That data from binary choices suffices to determine general choices is a non-trivial assumption, and one I drop. I identify a minimal choice consistency condition with an interesting ethical interpretation that must be rejected for Boonin-Vail's solution to be viable. By the result below, any such solution relies on path dependence. So one thing that is at stake in switching from a relation-theoretic to a choice-theoretic setting in an attempt to skirt impossibilities in population ethics is Path Independence. Similarly, choice-set-independence is the assumption that the permissibility of A when B is available is independent of the other options on the menu: if A is permissible in S and $B \in S$ is not, then B is not permissible in any menu that contains A. Some see failure of choice-set-independence of a theory as an obvious objection to it (Greaves, 2017; Mogensen, 2020). Meacham (2012) disagrees.[18] So do I. But the observation below establishes that, even if we assume something much weaker, impossibilities remain for population ethics.

Perhaps the most closely related work is Arrhenius's impossibility theorem for normative theories of population ethics (2004). Arrhenius's result appeals to a different set of assumptions than mine does, with no standard choice consistency conditions among the set. His ethical assumptions constrain choices between two options for *any* menu in which both appear. The ethical constraints stated here are only for certain binary choices, with Path Independence constraining choices in general. As Arrhenius concludes his essay, "We cannot exorcise the paradoxes of population ethics by giving up some formal condition like the transitivity of 'better than' or by rejecting consequentialism and switching to a normative framework. The paradoxes are a problem for any moral theory" (Arrhenius, 2004, p. 214). If we think of Path Independence as a minimal formal constraint on decision theory, one that substantially relaxes standard formal constraints in population ethics, a similar upshot is suggested by the present study.

[18] As constraints on choice functions as we have defined them, the two clauses of Meacham's IIA_d (2012, p. 273) are equivalent to each other. His constraint is one incarnation of the *Weak Axiom of Revealed Preference* (WARP) which is equivalent to weak order rationalizability when C has universal domain. As mentioned, weak order rationalizability is characterized by Properties α and β and is stronger than Path Independence (e.g., Moulin, 1985, p. 153).

While Mogensen (2020, p. 9) first states a principle—"binary independence"—in a way that sounds like choice-set-independence, the rest of the paper suggests his criticism concerns *binariness*. He objects to E-admissibility as a decision rule for imprecise probabilities on the grounds that it violates his condition, and argues for the Maximality rule. However, his example of a violation of binary independence for E-admissibility shows a failure of binariness, not just choice-set-independence. Moreover, Mogensen himself later produces an example of his preferred decision rule violating choice-set-independence (2020, p. 16).

3 Neutrality and Path Independence

Path Independence is a non-trivial generalization of rationalizability by either a weak order or even a preorder. As I pointed out above, while they have some nice features, path independent choice functions may not even be rationalizable by a binary relation. Though it allows for violations of choice-set-independence, Path Independence places certain limits on such violations. As I explain in this section, standard binary approaches to accommodating the important intuition of Neutrality by allowing for incompleteness retain Path Independence, but relax choice-set-independence. Relaxing choice-set-independence to Path Independence, then, opens up certain possibilities in population ethics. This section illustrates one such instance. While Path Independence is a weak assumption, as I show in the following section, a certain impossibility is robust to this generalization.

But first a possibility. A "basic intuition" about population ethics that many share is that bringing about additional (good) lives is of neutral value (Broome, 1994, p. 167). Narveson summarizes this intuition in the all-but-mandatory-to-cite slogan, "We are in favour of making people happy, but neutral about making happy people" (Narveson, 1973, p. 80).

> **Neutrality.** Suppose B is a profile generated by a single individual with welfare above the neutral level(s). Then, both A and $A \cup B$ are permissible in the binary choice between them.

The intuition underwriting Neutrality, claims Broome, "is deeply embedded in the way we think about the value of what we do. We generally simply ignore the effects of our actions on the world's population, even when the effects are predictable. This can only be because of the intuition that they are neutral in value. If that intuition turns out to be wrong, it will make a huge difference to the judgements we should make" (Broome, 2009a, p. 412). Yet, as Greaves observes in her review of the literature, "it turns out to be remarkably difficult to formulate any remotely acceptable axiology that captures this idea of neutrality" (2017, p. 8). In particular, despite finding Neutrality very compelling, Broome despairs of the possibility of incorporating it into an acceptable axiology (Broome, 2009b, p. 747).

If choice-set-independence—equivalently, rationalizability by a complete and transitive binary relation—is required, Neutrality would entail that A and $A \cup B$ are "equally good." A problem with this framing, as Broome points out, is that we could consider $A \cup B'$, where B', like B, is generated by a single individual (as Broome assumes, the same individual that generates B) enjoying a higher welfare level than in B (cf. 1994, p. 170). As an example, consider Parfit's procreation case again, where B is induced by a child with a mild disability and B' by that same child without that disability. According to Neutrality, A and $A \cup B'$ are also equally good. But if "equally good" is transitive—as Broome claims it to be "[a]s a matter of logic"—we have that $A \cup B$ and $A \cup B'$ are equally good, which strikes many as the wrong result.

Unlike choice-set-independence, Path Independence does not force Neutrality to render A and $A \cup B$ as equally good just because both are permissible. Suppose that C satisfies Path Independence and Neutrality. When A and $A \cup B$ are the only options available, both are permissible. The same goes for the case in which only A and $A \cup B'$ are available. If only $A \cup B$ and $A \cup B'$ are feasible, then only $A \cup B'$ is permissible. And when all three options are available, only alternatives A and $A \cup B'$ are permissible. There is no contradiction in these permissibility judgments so far as Path Independence goes.[19] The case from before of the couple deliberating about whether to have a child with a mild disability has this exact structure. If we assume that the permissibility of $A \cup B$ when A is available is independent of what other options are available, as choice-set-independence would have us do (Greaves, 2017, p. 10), we obviously do not get this result. Path Independence, on the other hand, allows for this form of context dependence.

It is well-known that there are binary relations that do not validate Broome's line of argument. As mentioned, discussions in population ethics often assume axiologies are preorders, allowing for forms of incompleteness (Thomas, 2018, p. 811). One way of developing Neutrality in a binary fashion is to count profiles induced by adding additional lives as incommensurable to the original profile. If A is incommensurable with both $A \cup B$ and $A \cup B'$, but $A \cup B'$ is better than $A \cup B$, there is no conflict with reflexivity or transitivity (cf. Rabinowicz, 2009). Choice functions generated by maximizing a (not necessarily complete) preorder satisfy Properties α, γ, and Aizerman's Axiom. That is, such choice functions are binary path independent choice functions (Moulin, 1985, Theorem 4), though they may fail to satisfy choice-set-independence.[20] So neither mere Path Independence nor binary Path Independence on their own generates Broome's contradiction. Still, even dropping the binariness assumption, Path Independence confronts substantive limitations in population ethics. This suggests that it is worth focusing directly on Path Independence.

4 The Persistence of the Mere Addition Paradox

We are now in a position to state the choice-theoretic versions of three important desiderata for theories of population ethics, and prove that the Mere Addition Paradox persists for path independent choice. The first is for a theory of population ethics to avoid Repugnant Conclusion.

> **Repugnant Conclusion.** For any profile A induced by population at a blissful welfare level, there is another profile generated by a population at a drab level, B, that is permissible in the choice between A and B.

Representative blissful and drab populations are depicted in Figure 1, where the width of a block indicates the size of the population and the height indicates the magnitude of welfare of the people in the population. Typically, the Repugnant Conclusion asserts that, for any blissful population, there is a drab population that is *better*. The Repugnant Conclusion as stated here in choice-theoretic terms is slightly weaker, not just because of the generalization from preference to choice, but also because it states only that the drab population is permissible in a binary choice—not that the blissful population is impermissible. As a result, the requirement that a theory avoid the Repugnant Conclusion is slightly stronger.[21]

[19] Such a path independent choice function C could be interpreted as selecting the set of best options according to $\{\succsim_1, \succsim_2\}$ such that $A \cup B' \succ_1 A \cup B \succ_1 A$ and $A \succ_2 A \cup B' \succ_2 A \cup B$.

[20] At the level of choice functions, maximizing a preorder R is indistinguishable from optimizing the extended complete and quasi-transitive relation $R' = R \cup \{(x, y) : \neg xRy \text{ and } \neg yRx\}$. The extended relation effectively converts incompleteness into indifference, and may fail to be transitive in its symmetric factor. Nevertheless, a crucial conceptual difference remains between R and R', between incompleteness and indifference.

[21] Parfit discusses similar weakened forms of the Repugnant Conclusion (e.g., 1984, p. 432). Ng appeals to an analogous weakening of the Repugnant Conclusion for binary relations in his trilemma: "the necessity to

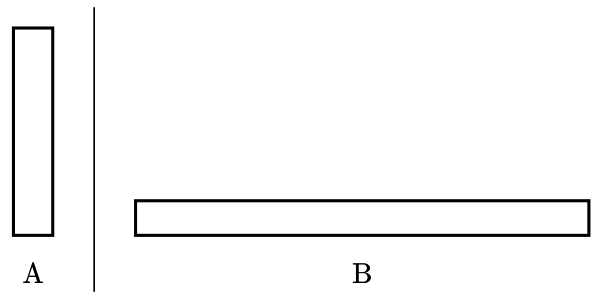

Figure 1: Blissful and Drab Populations

The next desideratum is "widely seen as (at least) the default hypothesis" (Thomas, 2018, p. 820, fn. 20) or even "*obviously* true" (Tännsjö, 2002, p. 357).[22] The mere addition of people with lives worth living to a population is better or at least not worse than the original population. Here I state the condition in choice-theoretic terms.

> **Mere Addition.** Suppose that A and B are profiles that contain only positive welfare levels. Then $A \cup B$ is permissible in the choice between A and $A \cup B$.

In Ng's presentation, Mere Addition states that $A \cup B$ is "better or at least not worse" than A (1989, pp. 237-238). Greaves likewise presents Mere Addition as asserting that $A \cup B$ is "not worse than" A (2017, p. 3). Here, I make only the weak assumption that $A \cup B$ is permissible in the binary choice between it and A.

Ng introduces the final desideratum which allows him to illustrate Parfit's Mere Addition Paradox in a more efficient way. It constrains choices between an unequal welfare distribution and an equal distribution with higher total or general welfare.

> **Non-Anti-Egalitarianism.** In a choice between a profile A with different welfare levels and a profile B of the same length, with a single welfare level, and with higher total welfare, A is not permissible.

Non-Anti-Egalitarianism, according to Ng, "is extremely reasonable. Total utility increases, average utility increases, the profile of utility distribution is more (and in fact perfectly) equal, and other things remain unchanged. There is simply no acceptable ground not to regard this as preferable" (1989, p. 238). According to Greaves, "violation of Non-Anti-Egalitarianism seems unacceptable" (2017, p. 6), and Thomas reports that "many people" find the principle "extremely compelling" (2018, p. 823).[23] While the generalized statement here does not assume a binary relation, it retains what so many find attractive about the original condition.

say that E is better than or *at least not worse than* A must still be regarded as an instance of the Repugnant Conclusion" (1989, p. 240, emphasis added). Similarly, Carlson writes, "I do not find [the weaked form of the Repugnant Conclusion] much less repugnant [...] and I believe that most people who object to [the Repugnant Conclusion] would share this judgment" (Carlson, 1998, p. 287).

[22] But see Kitcher (2000, p. 567).

[23] Thomas expresses a few reservations about the condition, including some skepticism about the presupposed notion of total welfare (2018, pp. 822-824). If I understand correctly, the suggested strategy to skirt some of these concerns (2018, p. 822-823, fn. 27) might be adapted for our purposes here, though the details will depend on how we think about total welfare.

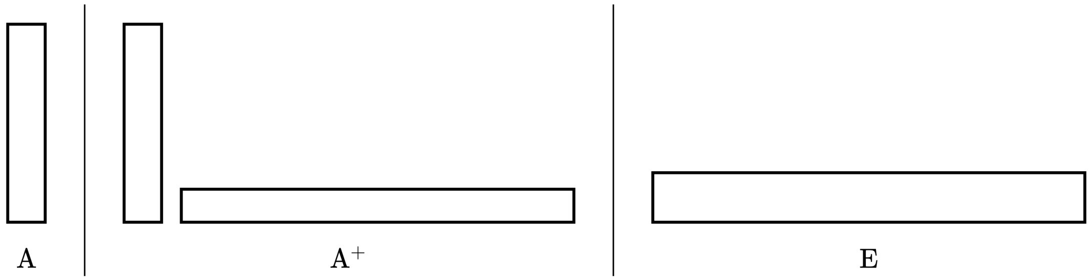

Figure 2: Three Alternatives

Observation. *Path Independence, Mere Addition, and Non-Anti-Egalitarianism together imply the Repugnant Conclusion.*

Proof. Suppose that C is path independent and that it satisfies both Mere Addition and Non-Anti-Egalitarianism. Fix a profile A as in the Repugnant Conclusion. Let $A^+ = A \cup B$, where B is a profile induced by a very large population with a uniform, drab level of welfare. Finally, let E be of the same length as A^+ with slightly higher total welfare, but uniformly distributed. (We assume that some such B and E can be found for sufficiently large populations to secure that the single welfare level in E is also drab.) These alternatives are depicted in Figure 2. By Mere Addition,

$$A^+ \in C(\{A, A^+\}). \tag{1}$$

By Non-Anti-Egalitarianism, we have

$$C(\{A^+, E\}) = \{E\}. \tag{2}$$

Now consider the three-way choice between A, A^+, and E. First, if $E \in C(\{A, A^+, E\})$, then, by Property α, $E \in C(\{A, E\})$, which gives us the Repugnant Conclusion and we're done. So assume for reductio that $E \notin C(\{A, A^+, E\})$. Next, suppose that $A^+ \in C(\{A, A^+, E\})$. By Property α, $A^+ \in C(\{A^+, E\})$, which is inconsistent with 2. So, $A^+ \notin C(\{A, A^+, E\})$. Finally, suppose that $A \in C(\{A, A^+, E\})$. Then, by Property α, $A \in C(\{A, A^+\})$. Since $C(\{A, A^+, E\}) = \{A\} \subseteq \{A, A^+\}$, by Aizerman's Axiom, $C(\{A, A^+\}) \subseteq C(\{A, A^+, E\})$. In particular, $A^+ \in C(\{A, A^+, E\})$, which is a contradiction. Since choice sets are nonempty, it must be the case that E is permissible in the three-way choice, and so, by Property α, E is permissible in the choice between it and A. Since A was arbitrary, the Repugnant Conclusion follows.[24] □

Notice that we made no assumptions about the composition of $\mathbb{A}$ beyond those involving, for any A, the existence of a suitable B to construct A^+ and a suitable profile E, and no assumptions about how C behaves beyond those explicitly mentioned. So no general inferences about what other properties C may have—such as binariness—are warranted.

[24] Using just the Path Independence axiom seems to me slightly less transparent. By Path Independence and 2, $C(\{A, A^+, E\}) = C(C(\{A^+, E\}) \cup C(\{A\})) = C(\{A, E\})$. Therefore, $E \in C(\{A, A^+, E\})$ if and only if $E \in C(\{A, E\})$. For reductio, suppose that $E \notin C(\{A, A^+, E\}) = C(\{A, E\})$. Then, since choice sets are nonempty, $C(\{A, A^+, E\}) = \{A\}$. Using Path Independence again, $C(\{A, A^+, E\}) = C(C(\{A, E\}) \cup C(\{A^+\})) = C(\{A, A^+\})$. By 1, $A^+ \in C(\{A, A^+\}) = C(\{A, A^+, E\})$, which is a contradiction. So, $E \in C(\{A, E\})$.

The choice-theoretic version of the Mere Addition Paradox is the conjunction of three initially plausible claims. For the options in Figure 2, (1) A^+ is permissible in the choice between it and A (by Mere Addition), (2) A^+ is not permissible in the choice between it and E (by Non-Anti-Egalitarianism), and (3) E is not permissible in the choice between it and A (by avoiding the Repugnant Conclusion). Unfortunately, given Path Independence, these three claims, and so the three principles that imply them, are inconsistent. If we are to avoid the paradox, either we have to surrender at least one of the substantive ethical desiderata—the negation of the Repugnant Conclusion, Mere Addition, or Non-Anti-Egalitarianism—or we have to accept that what is good (in terms of a theory of value) or permissible (in terms of a normative theory) depends, not just on what other options are available, but on the order of assessment as well.

5 Discussion

There are at least three broad types of reactions to the Mere Addition Paradox. The first is to accept the Repugnant Conclusion, which some argue for explicitly (e.g., Huemer, 2008). Even if this response is perceived as unattractive in many ways, the view that avoiding the Repugnant Conclusion is obviously and necessarily a desideratum on any acceptable account of population ethics seems to be declining in currency (Zuber et al., 2021).

A second response is to jettison one of the other substantive ethical desiderata. For instance, perfectionists reject Non-Anti-Egalitarianism. In the three-way choice depicted in Figure 2, perfectionism could yield that $C(\{A^+, E\}) = \{A^+\}$, contrary to 2. The idea is that having people who achieve "the best things in life"—"the best kinds of creative activity and aesthetic experience, the best relationships between different people, and the other things which do most to make life worth living"—is of sufficient value to override the gains in equality and total welfare that E makes over A^+ (Parfit, 1986, p. 161).

A third response, more relevant for my study here, is to modify the choice theory. Certain egalitarians, for example, may find considerations of equality compelling grounds to regard A^+ as impermissible in the presence of E in the three-option menu $\{A, A^+, E\}$, yet think that such considerations are insufficient to exclude A^+ from the choice set for the menu $\{A, A^+\}$. In the latter menu, invoking egalitarian considerations to exclude A^+ has the drastic effect of reducing the number of lives worth living. In addition to preserving the binary judgments in 1 and 2, this proposal goes, we have $C(\{A, A^+, E\}) = \{A\}$ since E effectively blocks A^+ even though E is not itself permissible. Johann Frick advocates precisely this choice pattern in a forthcoming paper (2021). Though he does not discuss Path Independence or Aizerman's Axiom, the relevant issue for my purposes is that this choice pattern violates Aizerman's Axiom and thus Path Independence: $\{A, A^+\} \subseteq \{A, A^+, E\}$, $C(\{A, A^+, E\}) \subseteq \{A, A^+\}$, but $C(\{A, A^+\}) \nsubseteq C(\{A, A^+, E\})$ since A^+ is in the former but not the latter choice set. Dropping Aizerman's Axiom would prevent the foregoing derivation of the Repugnant Conclusion that invokes Path Independence.

As I read Frick, what precise strategy he is suggesting is open to a couple of different interpretations. It could be, as a referee suggests, that the proposal is to maximize, for each menu S, a menu-relative binary relation R_S. As we will see, this is *substantially* weaker than appealing to a set of relations defined on the entire set $\mathbb{A}$ to rationalize choice as in the case of path independent choice functions. In fact, it is as weak as any constraint on choice could be. Or the proposal might be to individuate options more finely than they typically are in treatments of the Mere Addition Paradox. I will consider these interpretations in turn, and register some reservations about them. Then, I will explain what the contribution of this essay would be if this sort of objection to Path Independence were ultimately sustainable.

On the first interpretation, for each menu $S \subseteq \mathbb{A}$, there is a menu-relative transitive binary relation $R_S \subseteq S \times S$, and C selects the R_S-maximal elements of S. Speaking against this interpretation of Frick is that it would be somewhat obscure what binary relation he means to claim exhibits "non-transitivity *across* different option sets" (more on this below) (2021, p. 24). Nevertheless, on this proposal, the choice function induced by $\{R_S\}_{S\subseteq\mathbb{A}}$ by setting $C(S) = \{A \in S : \neg\exists B \in S \text{ such that } BP_SA\}$ violates choice-set-independence for the particular R_S that Frick stipulates in the context of the Mere Addition Paradox. It also violates Path Independence, which, again, prevents the derivation of the Repugnant Conclusion above. A philosophical concern about this interpretation of Frick's project is that, as far as rational choice theory goes, it may be something of a Pyrrhic victory for vindicating the ethical intuitions ingredient to the Mere Addition Paradox. Without further constraints, which Frick does not articulate, the resulting theory of choice is entirely vacuous, ruling out no choice behavior whatsoever. If we allow menu-relative relations to rationalize choice, in other words, we can rationalize *any* pattern of choices with some set of transitive menu-relative relations.[25] At the very least, there is a constructive challenge here to complete a seriously incomplete case. What principles relate the various menu-relative relations such that rational choice theory is not entirely trivialized?

Attempts to respond to an analogue of this trivialization challenge have been made for the second way of reading Frick's project. Rather than appealing to a set of menu-relative relations to rationalize a choice function, Frick may be suggesting that we rethink the way options are individuated in the context of the Mere Addition Paradox. This reading is strongly suggested by a number of claims he makes. For example, according to Frick, the presence of E "gives outcome A^+ a property in [the menu $\{A, A^+, E\}$]—that of being *unjust*—which it does not possess in the [menu $\{A, A^+\}$]," and "The fact that A^+ is unjust in the [three-option menu] makes it a *worse outcome*, all else equal, than A^+ when the only alternative is A" (Frick, 2021, p. 18). On this second interpretation of Frick's project, he is claiming that the standard ethical judgments in the two-option menus involved in the Mere Addition Paradox can be preserved and the Repugnant Conclusion avoided by distinguishing the option A^+ in certain binary choice contexts from the option A^+ in the menu $\{A, A^+, E\}$.

As with the first interpretation in terms of a set of menu-relative binary relations, option

[25] To see this, let C be any choice function on $\mathbb{A}$. For each menu S, define

$$R_S := \{(A, B) : A \in C(S) \text{ and } B \in S \setminus C(S)\}.$$

First, observe that R_S is (trivially) transitive. For any three options $A, B, C \in S$, if AR_SB, then $A \in C(S)$ and $B \in S \setminus C(S)$. It cannot be the case, then, that BR_SC since this would imply that $B \in C(S)$, which is a contradiction. So, the antecedent of transitivity—AR_SB and BR_SC—is vacuously satisfied. Consequently, R_S is transitive. Second, observe that the set of relations $\{R_S\}_{S\subseteq\mathbb{A}}$ rationalizes C. Suppose that $A \in C(S)$ for some arbitrary menu S. By construction, there is no $B \in S$ such that BP_SA since this would imply that $A \in S \setminus C(S)$, which is a contradiction. So, $C(S) \subseteq M(S, R_S)$. Conversely, since the R_S-maximal elements in S are just defined to be those in $C(S)$—the only options A such that, for some $B \in S$, BP_SA are ones in $S \setminus C(S)$—$M(S, R_S) \subseteq C(S)$. So *any* choice function can be rationalized by a set of menu-relative transitive binary relations. (A referee points out that this observation can be strengthened to apply to menu-relative weak orders. The relevant construction and proof are analogous.)

re-description is risky, threatening to deprive us of a substantive theory of rational choice. In the extreme case, complete menu-dependence of option individuation trivializes inter-menu consistency conditions. Such conditions, Sen observes, "have cutting power only when 'the same' alternative can be picked from two different sets—precisely what is ruled out by [extreme option individuation]" (1997, p. 754). Since, as far as I am aware, the worry about trivializing rational choice theory is more widely appreciated in the context of option re-description, I will just say that, although some, aware of the worry, have proposed principles governing option individuation (e.g., Broome, 1991, Sections 5.4–5.8), such principles remain very controversial (Buchak, 2013, Chapter 4; Baccelli and Mongin, MS). There are, furthermore, a couple of reasons to think that this second interpretation may not be what Frick intends in the end. First, he discusses the individuation strategy, not as one he is advocating, but as one a defender of choice-set-independence might be tempted to pursue in response to his alleged counterexample (albeit a strategy "not without merit" (2021, p. 17)). Second, and more tellingly, the choice pattern Frick advocates in the context of the Mere Addition Paradox would fail to witness a violation of choice-set-independence, which seems to be the target of much of his critical discussion. Once we distinguish A^+ in $\{A, A^+\}$ from $A^{+'}$ in $\{A, A^{+'}, E\}$, it is easy to see that $\{A, A^+\}$ is not a subset of $\{A, A^{+'}, E\}$. Hence, Property β (required for choice-set-independence) as well as Aizerman's Axiom (required for Path Independence) are vacuously satisfied.

One subtlety to attend to concerns what Frick has to say about foresight and avoiding money pumps. In response to the worry that the choice pattern he advocates in a case closely related to the one depicted in Figure 2 is susceptible to a money pump, Frick distinguishes scenarios in which it is known that certain options will be offered sequentially, and scenarios in which such knowledge is lacking. "The key point," he writes, "is that in a situation where I will, sequentially, be offered a choice between all three options, their ethical properties do not differ from those in the [three-option menu], in which all three options are *simultaneously* available to me" (2021, p. 24). Applied to the set of menus involved in the Mere Addition Paradox, Frick's foreknowledge condition and the ranking that he advocates in the three-option menu—$A \succ E \succ A^+$ (2021, p. 22)—imply that, when it is known that A, A^+, E will all be offered at some point, $C(\{A, A^+\}) = \{A\}$. As a result, foresight would mandate a violation of Mere Addition. In the presence of foresight—perhaps a natural assumption for applications of Path Independence as a constraint on order of *assessment*—the correct choice pattern, on Frick's view, would not violate Path Independence or choice-set-independence. In other words, a viable challenge to Path Independence, on a view like Frick's, would depend also on the absence of foreknowledge.

Suppose, however, that Frick does not have this second, option-individuation interpretation in mind, and that he is right about the appropriate choice pattern in the context of the Mere Addition Paradox.[26] Here, I will abstract from the worries about sets of menu-relative relations and the precise modification of rational choice he intends more generally. In this case, the analysis presented in this essay helps us to identify a crucial property at stake. Frick focuses on *context*-dependence, which is equivalent to choice-set-dependence. But, in fact, on his account, the dependence runs deeper. The relevant choices are even *path*-dependent. Put another way, the costs of Frick's approach to the Mere Addition Paradox in terms of ra-

[26]Frick's proposal depends on some further controversial assumptions, including the rejection of a form of consequentialism (2021, Section 6), that I will not address.

tional choice theory are higher than he reports; more must be given up. The choice-theoretic approach helps us to identify the relevant property—Aizerman's Axiom—and allows us to investigate attractive and precise ways of weakening the condition.

If we accept, at least for the sake of the argument, that Aizerman's Axiom must be relaxed, the extent to which it should be weakened is now an important question. The following property is a very natural one to consider.

> **Weak** γ. If A is uniquely permissible in menu S and permissible in menu T, then A is permissible in the menu $S \cup T$.

This property is interesting for at least five reasons. First, in the presence of Property α, Aizerman's Axiom implies Weak γ (Stewart, 2020b, Proposition 3). Since the converse does not hold, the conjunction of Property α and Weak γ generalizes Path Independence without simply dropping all of the content of Aizerman's Axiom completely. Second, the choice pattern that Frick advocates in the context of the Mere Addition Paradox is consistent with the conjunction of Property α and Weak γ. So the generalization is both substantive and relevant to our present concerns. Third, especially notable about the conjunction of Property α and Weak γ is that it offers a simple functional characterization of a generalization of pseudo-rationalizability called *weak pseudo-rationalizability.* If C satisfies Property α and Weak γ, then there exists a set $\{R_i\}_{i \in I}$ of acyclic binary relations (rather than weak, total, or preorders as in the case of pseudo-rationalizable choice functions) such that, for every menu S, $C(S) = \bigcup_{i \in I} M(S, R_i)$ (Stewart, 2020b, Theorem 4). We retain the possibility of describing determinations of permissibility in terms of maximizing a set of axiologies. Fourth, in addition to generalizing pseudo-rationalizability to sets of acyclic relations, weakly pseudo-rationalizable choice functions are the "multi-preference" or "multi-axiology" analogue of rationalizability simpliciter: rather than C's being rationalizable by a single acyclic binary relation, C is rationalizable by a set of such relations. (At least for finite menus, acyclicity is the property needed to guarantee that each relation induces its own choice function and so could serve as a general criterion of permissibility.) Weakly pseudo-rationalizable choice functions, in other words, take both binary choice functions and path independent choice functions as special cases. Fifth, weak pseudo-rationalizability allows us to avoid any appeal to menu-relative relations or re-describing alternatives to accommodate Frick's choice pattern.

Is weak pseudo-rationalizability the appropriate setting in which to work out a theory of population ethics? If Frick is right about the appropriate choice pattern in the Mere Addition Paradox, perhaps so. However, I doubt that this prospect, despite the five points of interest just adduced, will meet wide-spread enthusiasm. One crucial issue at stake in the move from Path Independence to weak pseudo-rationalizability is the transitivity of the members of the associated set of relations—of the legitimate axiologies—that rationalize choice. The "revealed" axiology in the choice pattern that Frick advocates is not transitive.[27] Compared to the ethical judgments in the Mere Addition Paradox, many have regarded transitivity as more sacrosanct. If, as in standard treatments, axiologies are not relativized to menus à la the first interpretation of Frick, Path Independence as a coherence constraint on choice

[27]The choice pattern that Frick defends in the context of the Mere Addition Paradox is consistent with maximizing a(n acyclic) binary relation: $R = \{(A, E), (A, A^+), (A^+, A), (E, A^+)\}$. Perhaps this is what he has in mind with the claim about "non-transitivity across different option sets"; although, as we have seen, since any choice function can be rationalized by a set of transitive menu-relative relations, there is no guarantee that a choice function so rationalized will be binary in this way.

immediately follows from maximizing a transitive axiology. Therefore, however compelling or sacrosanct transitivity is, Path Independence is at least as compelling or sacrosanct since it is strictly weaker. Put differently, Path Independence should not be a controversial requirement for most population ethicists. The relevant question is whether requiring it alone, rather than making the stronger assumption of a unique transitive axiology, allows for any further fruitful generality for population ethics. I have urged that it does. The analogous question arises for weak pseudo-rationalizability. In moving to sets of axiologies, Path Independence preserves the transitivity of the legitimate axiologies in the associated set. Relaxing Aizerman's Axiom to Weak γ, despite its several conceptual motivations and the possibility of avoiding the Mere Addition Paradox, surrenders the transitivity of the axiologies in the associated set. Defenders of transitivity, even if they are willing to move from a single axiology to a set of them, would likely balk. The approach undertaken in this essay allows us to see an important connection between the classic debate about the status of transitivity in population ethics and the plausibility of different coherence constraints for a more general choice-theoretic framework.

References

Aizerman, M. and A. Malishevski (1981). General theory of best variants choice: Some aspects. *Automatic Control, IEEE Transactions on 26*(5), 1030–1040.

Aizerman, M. A. (1985). New problems in the general choice theory. *Social Choice and Welfare 2*(4), 235–282.

Arrhenius, G. (2004). The paradoxes of future generations and normative theory. In T. Tännsjö and J. Ryberg (Eds.), *The Repugnant Conclusion*, Volume 15, pp. 201–218. Dordrecht: Springer.

Arrow, K. J. (1963). *Social Choice and Individual Values* (2 ed.). New York: Wiley.

Baccelli, J. and P. Mongin (MS). Can redescriptions of outcomes salvage the axioms of decision theory? Unpublished Manuscript.

Boonin-Vail, D. (1996). Don't stop thinking about tomorrow: Two paradoxes about duties to future generations. *Philosophy & Public Affairs 25*(4), 267–307.

Bossert, W., M. J. Ryan, and A. Slinko (2009). Orders on subsets rationalised by abstract convex geometries. *Order 26*(3), 237–244.

Bossert, W. and K. Suzumura (2009). External norms and rationality of choice. *Economics and Philosophy 25*(02), 139–152.

Broome, J. (1994). The value of a person. *Proceedings of the Aristotelian Society, Supplementary Volumes 68*, 167–185.

Broome, J. (1995, Originally Published in 1991). *Weighing Goods: Equality, Uncertainty and Time.* Oxford: Basil Blackwell.

Broome, J. (2009a). Reply to Rabinowicz. *Philosophical Issues 19*, 412–417.

Broome, J. (2009b). Reply to Vallentyne. *Philosophy and Phenomenological Research 78*(3), 747–752.

Buchak, L. (2013). *Risk and Rationality.* Oxford: Oxford University Press.

Carlson, E. (1998). Mere addition and two trilemmas of population ethics. *Economics & Philosophy 14*(2), 283–306.

Chambers, C. P. and F. Echenique (2016). *Revealed Preference Theory*, Volume 56. Cambridge University Press.

Chambers, C. P. and M. B. Yenmez (2017). Choice and matching. *American Economic Journal: Microeconomics 9*(3), 126–147.

Danilov, V. and G. Koshevoy (2005). Mathematics of Plott choice functions. *Mathematical Social Sciences 49*(3), 245–272.

Dubra, J., F. Maccheroni, and E. A. Ok (2004). Expected utility theory without the completeness axiom. *Journal of Economic Theory 115*(1), 118–133.

Fishburn, P. C. (1973). *The Theory of Social Choice*, Volume 264. Princeton University Press Princeton.

Frick, J. (Forthcoming, 2021). Context-dependent betterness and the mere addition paradox. In J. McMahan, T. Campbell, J. Goodrich, and K. Ramakrishnan (Eds.), *Ethics and Existence: The Legacy of Derek Parfit*, pp. 1–27. Oxford: Oxford University Press.

Greaves, H. (2017). Population axiology. *Philosophy Compass 12*(11), 1–15.

Greaves, H. and T. Ord (2017). Moral uncertainty about population axiology. *Journal of Ethics and Social Philosophy 12*(2), 135–167.

Helzner, J. (MS). Admissibility in a logical framework. Manuscript.

Huemer, M. (2008). In defence of repugnance. *Mind 117*(468), 899–933.

Johnson, M. R. and R. A. Dean (2001). Locally complete path independent choice functions and their lattices. *Mathematical Social Sciences 42*(1), 53–87.

Kitcher, P. (2000). Parfit's puzzle. *Noûs 34*(4), 550–577.

Koshevoy, G. A. (1999). Choice functions and abstract convex geometries. *Mathematical Social Sciences 38*(1), 35–44.

Levi, I. (1980). *The Enterprise of Knowledge.* MIT Press, Cambridge, MA.

Levi, I. (1986). *Hard choices: Decision making under unresolved conflict.* Cambridge University Press.

Levi, I. (2008). Convexity and separability in representing consensus. In K. Basu and R. Kanbur (Eds.), *Arguments for a Better World: Essays in Honor of Amartya Sen*, Volume 1: Ethics, Welfare, and Measurement, Chapter 11, pp. 193–212. Oxford University Press.

Meacham, C. J. (2012). Person-affecting views and saturating counterpart relations. *Philosophical Studies 158*(2), 257–287.

Mogensen, A. L. (2020). Maximal cluelessness. *The Philosophical Quarterly 71*(1), 141–162.

Monjardet, B. and V. Raderanirina (2004). Lattices of choice functions and consensus problems. *Social Choice and Welfare 23*(3), 349–382.

Moulin, H. (1985). Choice functions over a finite set: a summary. *Social Choice and Welfare 2*(2), 147–160.

Narens, L. and B. Skyrms (2020). *The Pursuit of Happiness: Philosophical and Psychological Foundations of Utility.* Oxford: Oxford University Press.

Narveson, J. (1973). Moral problems of population. *The Monist 57*(1), 62–86.

Nehring, K. (1997). Rational choice and revealed preference without binariness. *Social Choice and Welfare 14*(3), 403–425.

Ng, Y.-K. (1989). What should we do about future generations? Impossibility of Parfit's theory x. *Economics & Philosophy 5*(2), 235–253.

Parfit, D. (1984). *Reasons and Persons.* Oxford: Oxford University Press.

Parfit, D. (1986). Overpopulation and the quality of life. In P. Singer (Ed.), *Applied Ethics*, pp. 145–164. New York: Oxford University Press.

Parfit, D. (1996). Acts and outcomes: A reply to Boonin-Vail. *Philosophy & Public Affairs 25*(4), 308–317.

Pedersen, A. P. (2009). Pseudo-rationalizability over infinite choice spaces. Technical report, Carnegie Mellon University.

Plott, C. R. (1973). Path independence, rationality, and social choice. *Econometrica: Journal of the Econometric Society 41*(6), 1075–1091.

Rabinowicz, W. (2009). Broome and the intuition of neutrality. *Philosophical Issues 19*(Metaethics), 389–411.

Savage, L. (1972, originally published in 1954). *The Foundations of Statistics.* New York: John Wiley and Sons.

Sen, A. (1977). Social choice theory: A re-examination. *Econometrica: Journal of the Econometric Society 45*(1), 53–88.

Sen, A. (1985). Well-being, agency and freedom: The Dewey lectures 1984. *The Journal of Philosophy 82*(4), 169–221.

Sen, A. (1994). Non-binary choice and preference:: A tribute to Stig Kanger. In D. Prawitz, B. Skyrms, and D. Westerstahl (Eds.), *Proceedings of the Ninth International Congress on Logic, Methodology and Philosophy of Science*, Volume 134 of *Studies in Logic and the Foundations of Mathematics*, Uppsala, Sweden, pp. 913–924. International Congress on Logic, Methodology and Philosophy of Science: Elsevier Science B.V.

Sen, A. (1997). Maximization and the act of choice. *Econometrica: Journal of the Econometric Society 65*(4), 745–779.

Sen, A. (2004). *Rationality and Freedom.* Harvard University Press.

Stewart, R. T. (2016). Conditional choice with a vacuous second tier. *Synthese 193*(1), 219–243.

Stewart, R. T. (2020a). A hyper-relation characterization of weak pseudo-rationalizability. *Journal of Mathematical Psychology 99*, 102439.

Stewart, R. T. (2020b). Weak pseudo-rationalizability. *Mathematical Social Sciences 104*, 23–28.

Suzumura, K. (1983). *Rational choice, Collective decisions, and Social Welfare.* Cambridge [etc.]: Cambridge University Press.

Tännsjö, T. (2002). Why we ought to accept the repugnant conclusion. *Utilitas 14*(3), 339–359.

Temkin, L. S. (1987). Intransitivity and the mere addition paradox. *Philosophy & Public Affairs 16*(2), 138–187.

Thomas, T. (2018). Some possibilities in population axiology. *Mind 127*(507), 807–832.

Von Neumann, J. and O. Morgenstern (2007, originally published in 1944). *Theory of Games and Economic Behavior (Commemorative Edition)*. Princeton University Press.

Zuber, S., N. Venkatesh, T. Tännsjö, C. Tarsney, H. O. Stefánsson, K. Steele, D. Spears, J. Sebo, M. Pivato, T. Ord, et al. (2021). What should we agree on about the repugnant conclusion? *Utilitas*, 1–5.