

The Monstrous Conclusion

This is a pre-print version. Please cite the final version, forthcoming in Synthese.

Luca Stroppa
University of St Andrews
Ls330@st-andrews.ac.uk

Abstract

This paper introduces the Monstrous Conclusion, according to which, for any population, there is a better population consisting of just one individual (the Monster). The Monstrous Conclusion is deeply counterintuitive. I defend a version of Prioritarianism as a particularly promising population axiology that does not imply the Monstrous Conclusion. According to this version of Prioritarianism, which I call Asymptotic Prioritarianism, there is diminishing marginal moral importance of individual welfare that can get close to, but never quite reach, some upper limit. I argue that Asymptotic Prioritarianism faces a theoretical cost, that I call the Absolute Priority Principle. However, the Absolute Priority Principle is an extreme version of what I call the Trade-off Condition, an already noteworthy problem facing other (more widely endorsed) versions of Prioritarianism. I conclude that it is better for a theory to imply the Absolute Priority Principle and avoid the Monstrous Conclusion than to imply the Monstrous Conclusion and the Trade-off Condition. The potential for Asymptotic Prioritarianism is substantial.

Keywords: Repugnant Conclusion, Utility Monster, Impossibility Theorems, Prioritarianism, Population Axiology

1. Introduction

How much should we spend now to prevent the long term harms of climate change? Should we prioritize spending on current healthcare versus preparing for future pandemics? What should we pay to mitigate the future risks of nuclear waste? These decisions affect both who is born, how many people are ever born, and how well off they will be. For these decisions, we need variable-population ethics, or ‘population ethics’ for short.

A key part of population ethics is population axiology, which is concerned with evaluating what makes one population better than another. According to pure consequentialists, whether one population is better than another is all that matters in choosing between them. Non-consequentialists disagree, maintaining that non-axiological considerations such as rights and agent-relative prerogatives matter. However, most non-consequentialists believe that axiology can also guide us in deciding which population to choose, for example, when no rights are violated and there are no agent-relative considerations at stake.

The betterness ranking of populations depends at least in part on how well off the individuals are in these populations, whatever well-being consists in.¹ For example, it is very intuitive that, if two populations share the same individuals, but every individual is better off in one population than in the other, then the population in which each person is better off is better overall.

It is notoriously challenging to find a plausible population axiology. Indeed, several impossibility theorems show that any theory of population axiology will have some disturbing implication or another ((Arrhenius, 2000; Carlson, 1998; Kitcher, 2000); See (Greaves, 2017) for an overview).

This paper introduces the Monstrous Conclusion, a new serious challenge for population axiologies. Consider the following:

The Monstrous Conclusion: for any population, there is a better population consisting of one individual.

The Monstrous Conclusion is deeply counterintuitive. I argue that we should concentrate our attention on theories that can avoid it.

The Monstrous Conclusion bears a striking resemblance to Nozick's Utility Monster objection (1974, p. 41), one of the most widely cited objections to utilitarianism (Briggs & Nolan, 2015; Kamm, 2015; Miller, 2021; Rosenqvist, 2020;

¹ There are three families of theories about what it may mean for someone to be 'better off'. On 'mental state theories', how good a life is depends on the duration and intensity of pleasurable mental states over painful mental states. On 'preference satisfaction theories', how good a life is depends on the preferences and desires of the person who lives that life. On 'Objective list theories', how good a life is depends on whether it includes certain objective goods. For an influential summary of these three families of theories, see (Parfit, 1984, pp. 493–502).

Vallentyne, 1991). According to Nozick, ‘Utilitarian theory is embarrassed by the possibility of utility monsters who get enormously greater gains in utility from any sacrifice of others than these others lose’ (Nozick, 1974, p. 41). Nozick’s case concerns a fixed population. The two possible populations are the status quo population and one in which everyone has their resources transferred to a ‘Monster’ and suffers the result. By contrast, the Monstrous Conclusion I am proposing does not require the same people to exist in both populations: in one possible population there are, say, billions of very happy people, while in the other, there’s just the Monster. While Nozick’s Utility Monster is extremely influential, its variable-population analogue is underexplored: before we delve into the Monstrous Conclusion, it would make sense to examine why.

The Utility Monster has been overshadowed by the Repugnant Conclusion in population ethics, in part because of how Parfit defended that the Repugnant Conclusion had independent interest despite Nozick’s pre-existing Utility Monster. Parfit argues that Nozick’s Utility Monster is possible only with major changes to the laws of nature, that it “is a deep impossibility” (1986, p. 389), and that we cannot successfully imagine it. Specifically, he asserts that the welfare of the Monster needs to be higher than what the laws of nature allow, and therefore “[i]t seems a fair reply that we cannot imagine, even in the dimmest way, the life of this Utility Monster. And this casts doubt on the force of the example” (1986, p. 389). However, it is not at all clear why the laws of nature should prevent someone from having an arbitrarily great amount of welfare.

Instead, according to Parfit, the Repugnant Conclusion is not a deep impossibility. According to the Repugnant Conclusion, for any population, there is some better population consisting only of lives that are barely worth living.² Parfit argues that “the difference [between the Repugnant Conclusion and the Utility Monster] is that the greater sum of happiness comes from a vast increase, not in the quality of one person’s life, but in the number of lives lived. And [the Repugnant Conclusion] is neither deeply impossible, nor something that we cannot imagine. [...] So the example cannot be questioned as one that we can hardly understand”. (1986, p. 389)

² This is a modification of the formulation in (Parfit 1986, p. 388). I assume here and throughout the paper that things other than welfare and number of people are equal across the populations I compare. For example, I assume that the two populations do not differ in desert, autonomy, genesis, aesthetic value, and so on.

By contrast, I contend that both the Repugnant Conclusion and the Monstrous Conclusion are similarly important for a satisfactory population axiology to avoid. This is because I find Parfit's case for the Repugnant Conclusion being more important than Nozick's Utility Monster as a test for moral theories unconvincing. What I find unconvincing is not that the welfare of Nozick's Utility Monster may be hard to imagine. Rather, I find unconvincing that, as there are arguments for the unimaginability of Nozick's Utility Monsters, there are also arguments for the unimaginability of the Repugnant Conclusion of comparable strength. Indeed, one of the other rare uses of Nozick's Utility Monster in variable population context³ has been to argue that we cannot imagine the enormous quantity of people necessary for the Repugnant Conclusion ((Ng, 1989, p. 242); for further discussion of the unimaginability of large numbers, see (Broome, 2004, p. 57; Gustafsson, 2022; Huemer, 2008, p. 904; Tännsjö, 2002)).

In response, while arguably we may not be able to grasp *how many* people are involved, we can grasp that *no number of people* would suffice (Mogensen, 2022; Parfit, 2016, p. 111; Pummer, 2013; Temkin, 2012, pp. 35, 121–122, 155). Similarly, while we may not be able to grasp *how good* the Monster's life is, we can still grasp that *there is no* amount of goodness for the Monster that would suffice. Moreover, even Parfit says that Nozick's Utility Monster 'may provide a partial test for our moral principles. We cannot simply ignore imagined cases' (1986, 389).

Meanwhile, others have argued for views where the life of the Utility Monster could also be rather mundane, for example on views where there is non-diminishing value to experiencing a given momentary quality of life for a longer time (Arrhenius, 2000, p. 44; McTaggart & Broad, 1968; Parfit, 2004, 2012). Moreover, even if the Repugnant Conclusion were compatible with the laws of nature but the Monstrous Conclusion was not, avoiding one but not the other would make our population axiology contingent on these facts in implausible ways (Arrhenius, 2000, pp. 50–51).

While enormous efforts have gone to exploring ways of avoiding the Repugnant Conclusion, there has been little discussion on how to avoid the Monstrous Conclusion (see (Blackorby et al., 2005; Ng, 1989; Parfit, 2016; Temkin, 2012; Thomas, 2018) and many others). Given how many ways there are that the Repugnant Conclusion and Monstrous Conclusion would be of comparable importance, there is clearly an

³ Others are (Pivato, 2014, 2018). They will be discussed in section 2.2.

imbalance between attempts to address the two. In this paper, I begin to tackle this imbalance.

I argue that a promising way to avoid the Monstrous Conclusion is a version of Prioritarianism. Specifically, after presenting the Monstrous Conclusion in greater detail, I argue that a promising way to avoid it is via a hitherto undefended version of Prioritarianism which I call Asymptotic Prioritarianism. On Asymptotic Prioritarianism, while there is always moral importance to increasing an individual's welfare, but there is some limit, or asymptote, to the moral importance of one individual's welfare. On the other hand, there is no asymptote to the moral importance of multiple other people's welfare.

Nevertheless, every population axiology has its counterintuitive consequences, Asymptotic Prioritarianism included. Therefore, to establish what counterintuitive consequences the 'least-bad' population axiology must compromise on, more must be done to develop axiologies that can avoid the Monstrous Conclusion.

2. The Monstrous Conclusion

2.1 The Monster

According to the Monstrous Conclusion, for any population, there is a better population consisting of just one individual who is sufficiently better off (the Monster). The Monstrous Conclusion is deeply counterintuitive.

To appreciate the implausibility of this claim, consider that the Monstrous Conclusion applies to *any* population. Suppose we have a vast population, say a hundred billion people, each leading an excellent life. The lives could be as good as you like, much better than anyone has today. They enjoy great pleasures, have deep understanding, fulfill challenging projects, and develop meaningful relationships. They also do not experience pain or agony, there is no evil in the world and no-one is subject to malice or other's domination. Nevertheless, the Monstrous Conclusion implies that there is a *better* population consisting of a single individual, the Monster. I expect most people will find the Monstrous Conclusion deeply implausible, even impossible to believe.

While the Monstrous Conclusion has similarities with Nozick's Utility Monster objection, the Monstrous Conclusion enables us to more clearly identify the problems with giving the Monster too much moral importance (1974, p. 41). The key difference

between the Monstrous Conclusion and Nozick's Utility Monster objection is that Nozick considers a fixed population case, whereas the Monstrous Conclusion is about variable populations. This difference has a number of important upshots.

Firstly, a number of negative features that one can appeal to in order to explain why all the resources should not be given to the Utility Monster in Nozick's objection are not available to resist the Monstrous Conclusion. For example, the rest of the population is harmed in Nozick's objection but not in the Monstrous Conclusion (Broome, 2004; Bykvist, 2007; Heyd, 1988; Parfit, 1984, pp. 487–489) – the Monster just exists by themselves, and no one needs to be harmed. Similarly, Nozick's case involves vast inequality between the Utility Monster and the rest of the population, whereas the Utility Monster exists alone in the Monstrous Conclusions case, so cannot involve any inequality (Nebel, 2017, p. 898; Otsuka, 2012, p. 370; Voorhoeve & Fleurbaey, 2016). As the Monstrous Conclusion is highly implausible despite lacking these bad features, there must be sufficient explanation for the Monstrous Conclusion's implausibility that does not rely on these features.

Secondly, none of the standard positive features that philosophers appeal to in order to rank populations straightforwardly count against the Monstrous Conclusion.⁴ The population only containing the Monster has: higher total welfare (Arrhenius, 2000, pp. 37–51; Huemer, 2008; Parfit, 1984, pp. 397–389; Tännsjö, 2002); average welfare (Grill, 2023; Parfit, 1984, p. 387; Pressman, 2015); no inequality (Nebel, 2017, p. 898; Otsuka, 2012, p. 370; Voorhoeve & Fleurbaey, 2016); there are more perfectionist goods (Beard, 2020; Parfit, 2016) and so on...

As the Monstrous Conclusion is deeply counter-intuitive, there must be more to the value of populations than the literature has assumed so far. Given that the Monstrous Conclusion is deeply counter-intuitive despite these differences, the Monstrous Conclusion is an excellent tool for understanding what is intrinsically unsatisfactory about the Utility Monster population.

2.2 Most theories imply the Monstrous Conclusion

Having motivated the independent interest in the Monstrous Conclusion, in this subsection I explore how extant population axiologies fare with respect to avoiding the

⁴ This will be explained in greater detail in section 2.2.

Monstrous Conclusion. After showing the limits of a seemingly simple solution for avoiding the Monstrous Conclusion, in this sub-section I show that the Monstrous Conclusion is implied by a range of major population axiologies. For reasons of space, I cannot go through all population axiologies defended in the literature. I therefore restrict my analysis to the Total View, the Average View, views combining Total and Average aspects, Critical Level and Critical Range Views, and Person Affecting Views. These are the views covered in a recent influential survey (Greaves 2017). I conclude this section by pointing out two views defended in the literature that do avoid the Monster, but do so in undesirable ways.

A seemingly simple, but ultimately unsatisfactory, way to escape the Monster is with an assumption of bounded utility. On the assumption of bounded utility, there is an upper limit to how much welfare any being can have. If there is a limit to how much welfare any being can have, then this also limits how much the Monster can contribute to the value of a population – I call this the “contributory value” of the Monster to a population. According to the Monstrous Conclusion, for any population, there is a better population consisting of only one individual with greater contributory value than the given population. If welfare is bounded above, and if, for some given population, the Monster needs to have a contributory value greater than the limit to be better than the given population, then the assumption of bounded utility prevents the Monstrous Conclusion in *all* theories, including the Total and the Average View.

Bounded utility is often used as a convenient assumption in economics. The primary justification given for bounded utility is that it is the easiest way to avoid certain ‘paradoxes’ that have unbounded payoffs, most famously the St Petersburg and Pasadena paradoxes, but adopting bounded utility is not the only solution to such paradoxes (Arrow, 1971, p. 92; Cowen & High, 1988; Nover & Hájek, 2004). See (Blackorby et al., 2005, p. 91; Kreps, 2013, pp. 11–13; Savage, 1954; Von Neumann & Morgenstern, 1944) for those who adopt bounded utility and see (Buchak, 2017; Dillenberger & Vijay Krishna, 2014; Fishburn, 1976; Hájek & Smithson, 2012; Kosonen, 2022) for those who do not, among many others. As both bounded and unbounded utility are widely used in economics depending on the application of a given theory or model, and given that population ethics is meant to be applied to any logically possible population (Arrhenius, 2000, pp. 31–35, Forthcoming, p. 41; Beckstead & Thomas, 2023; Gustafsson, 2020; Huemer, 2008; Thomas, 2018), population ethics should not merely assume bounded utility. Indeed, from the outset of the field of

population ethics, unbounded utilities were accepted as at least a partial test of our theories (see Parfit, 1984, p. 389).

Potentially more importantly, some philosophical debates in prudential axiology have implications for whether welfare must have some upper limit. As an example of these debates, consider the debate between those who argued that increasing the duration of a happy life cannot improve this life beyond a certain limit (Beglin, 2017; Kagan, 2012; Smuts, 2011; Temkin, 2008; Williams, 1973, pp. 224–232) and those who have argued the opposite (Beckstead & Thomas, 2023, pp. 13–14; Bruckner, 2012; Fischer & Mitchell-Yellin, 2014; Gorman, 2017; Greene, 2017). I am personally unconvinced that any such arguments have established that individual welfare must have an upper limit. Nevertheless, I do not argue for that here, nor do I need to in order to assume unbounded utility.

Indeed, assuming unbounded utility has methodological advantages regardless of whether utility is actually bounded or unbounded. If utility is unbounded, then we have an account of population ethics that captures the full range of cases. If utility is bounded, then we still have a theory covering the full range of cases by considering the part of our theory up to that bound, but we can also determine which important features in population ethics depend on these controversial debates in prudential axiology. Therefore, whether or not it is the case that utility is in fact bounded, significant results can be obtained by using an assumption of unbounded utility. Thus, I assume unbounded utility for the rest of this paper. Let us now analyse how some views imply the Monstrous Conclusion.

The Total View. According to the Total View, one population is better than another if the total amount of welfare is higher (Parfit 1984, 387-389; Arrhenius 2000, 37-51; Huemer 2008; Tännsjö 2004). The Total View implies the Monstrous Conclusion. For any population, that population will have a given amount of total welfare. So long as the Monster has more welfare than that, the population containing the Monster only is better according to the Total View. Therefore, for any population, there is a better population consisting of a single individual (the Monster): this is the Monstrous Conclusion.

The Average View. According to the Average View, one population is better than another if the average amount of welfare of each individual is higher (Parfit 1984, 387; Pressman 2015; Grill 2023). The Average View implies the Monstrous Conclusion. For any population, that population has some average welfare level. As the

Monster is the only member of the single individual population, its welfare is the same as the average welfare of the population. So long as the Monster has more welfare than the average welfare of the given population, the population containing only the Monster is better according to the Average View. Therefore, for any population, there is a population consisting of a single individual that is better. This is the Monstrous Conclusion.

Combined Views are a family of theories that aims to combine the Total View and the Average View, in such a way that the strengths of each view cover the other view's weaknesses. Typically, these theories tend towards the Total View when evaluating small populations, and towards the Average View when evaluating larger populations (Hurka, 1983; Ng, 1989; Sider, 1991).⁵ Since the Monstrous Conclusion is a weakness of both the Total View and the Average View, Combined Views inherit this problem: for any population, there is a Monstrous Population with greater total and average welfare.

Critical Level and *Critical Range Views* are modifications of the Total View aimed at preventing populations containing only barely worth living lives having unbounded value. These views either adopt a sufficiently good level, where lives that are not sufficiently good count against a population (these are Critical Level Views: see (Blackorby et al., 1997, 1998, 2005)) or an extended range around the neutral level that does not count in favour of a population (these are Critical Range Views: see (Gustafsson, 2020; Rabinowicz, 2009, 2022; Thornley, 2022)).

However, as these views only adjust the contributory value of welfare around the neutral level, they do nothing to prevent the welfare of one individual having unbounded contributory value. Therefore, for any population, that population will have a given total contributory value adjusted by the critical level or range. So long as the Monster has more contributory value than that (adjusted by the critical level or range), the population with only the Monster in it is better according to Critical Level and Critical Range Views. Thus, for any population, there is a better population consisting of a single individual (the Monster): this is the Monstrous Conclusion.

⁵ There are important differences between Sider's theory and the theories by Hurka and Ng. On Hurka and Ng, the contributive value of additional people of a population depends on the average welfare of the population. Sider's theory resembles a theory of marginal values, where the contributive value of additional lives is less the more people there are with a welfare higher than the one of the additional people.

Person Affecting Views are a family of views according to which welfare that involves making people better and worse off has a distinctively central moral importance compared to adding to a population of people who have positive welfare (Heyd, 1988; Roberts, 2009; Ross, 2015; Temkin, 2012, sec. 12).⁶ *Person Affecting Views* imply the Monstrous Conclusion. As the Monstrous Conclusion just requires that there is *some* one-person population which is better, this individual may be present in the original population.

To see why *Person Affecting Views* imply the Monstrous Conclusion, let us suppose that the individual is one of those who exists in the given population. Now suppose that that person gets much more welfare than the whole of the given population put together. As that person is much better off, there is enormous comparative benefit to producing the one individual population, as well as much greater welfare overall. Thus, as there is both greater comparative benefit and welfare overall, according to *Person Affecting Views* the one individual population is better: this is the Monstrous Conclusion.

To diagnose the implausibility of the Monstrous Conclusions we must look beyond these major population axiologies. Specifically, I will look into two kinds of explanations as to why populations can be better than the Monster. I will reject one of the explanations, while I will incorporate the other in my proposed theory in the next section. I begin by considering and rejecting *Sufficientarianism* before the axiology which explicitly aims to avoid the Monster, Pivato's *Rank-additive Population Axiology* (2018).⁷ I will also highlight the connections between Pivato's work and the version of Prioritarianism I will defend in the next section.

Sufficientarianism is the family of views which claim that absolute priority should be given to the welfare of people below a certain welfare threshold (Bossert et al., 2022; Brown, 2005; Casal, 2007; Crisp, 2003; Frankfurt, 1987; Hirose, 2016). *Sufficientarianism* captures the attractive idea that it is more important to make people sufficiently well off than to make the already well off even better off, albeit in an extreme form. This idea can be used to resist the Monstrous Conclusion by arguing that

⁶ See (Arrhenius 2000, 114-138; Greaves 2017, 11-16) for more information and some important distinctions on *Person Affecting views*.

⁷ Pivato mentions the Utility Monster in variable population context also in (2014), where he points out that population axiologies based on adding individual wellbeing imply that "it is better to starve N moderately happy people, just so that one person can achieve the 'Nirvana' state y." (Pivato, 2014, p. 38).

there is absolute priority to creating people who are sufficiently well off, rather than making the Monster better and better off.

For example, on the sufficientarian “head-count approach”, the aim is to ensure that “as many people as possible have enough” (Frankfurt 1987 p. 31). That is, populations are ranked according to the number of people who have enough, therefore, as the Monstrous population only has one person who has enough, any population with at least two people who have enough is better than the Monstrous population.

There is a sense in which avoiding the Monstrous Conclusion necessarily involves an ‘absolute’ priority claim – the welfare of sufficiently many sufficiently well off people has priority over the welfare of one individual, no matter how much welfare is at stake. Nevertheless, I shall argue that Sufficientarianism faces fatal defects that can be avoided by capturing this ‘absolute’ priority claim in the form of Prioritarianism instead.

In particular, Sufficientarian theories face a serious, well-known problem. These theories rely on a sufficiency threshold below which any welfare increase gets absolute priority over welfare increases above the threshold: this creates a particularly extreme kind of discontinuity in the welfare spectrum. More precisely, Sufficientarianism implies the following. Consider two very close welfare levels, call them w_1 and w_2 . w_1 is just below the threshold, and w_2 just above it. Since, on Sufficientarianism, absolute priority should be given to the welfare of people below the threshold, a small improvement to an individual at w_1 is more important than any improvement to w_2 , no matter how great the improvement to w_2 is. Since w_1 and w_2 can be arbitrarily similar, it is hard to believe that we should treat them so dissimilarly.

As there is no plausible candidate for a point in the welfare spectrum that demarcates such a stark difference, Sufficientarianism’s requirement of this discontinuity in the welfare spectrum is widely considered implausible (Arneson, 2000, p. 56, 2002, p. 194; Casal, 2007, p. 317; Dorsey, 2014, pp. 50–53; Holtug, 2010, pp. 207, 227–31; Timmer, 2022, pp. 308–309).

Finally, a more sophisticated approach is developed by Pivato in his paper examining *Rank-additive theories* of population axiology (2020), where one of his explicit aims is to respect a “no utility monsters” adequacy condition for any theory of population ethics. Rank-additive theories “admit an additively separable representation”, meaning that the overall value of a population can be represented by the contributory value of each individual in the population added together, like the

“classical utilitarian” or “prioritarian” value functions (Pivato 2020 p.863). What makes Rank-additive theories distinctive compared to these other additively separable theories is that “people are ranked in order from lowest to highest lifetime utility, and different transformations can be applied to different entries in this ranking” (Pivato 2020 p.863). In other words, the moral significance of a given welfare increase at a given welfare level (the transformation of welfare increase into contributory value increase) could be different depending on whether that person is the 3rd best off or the 10,000th best off in their population.

On the one hand, Pivato is correct to be concerned with avoiding Nozick’s Utility Monster and to adopt a population ethics analogue of Nozick’s classic fixed population case. However, Pivato’s “No utility monsters” axiom is too strong to be considered an adequacy condition for any population axiology. Pivato says that:

“The [No utility monsters] axiom rules out Nozick’s (1974) Utility Monster paradox. It says that for *any* finite population size N , there exists a [population] (presumably involving a larger number of people) which is better than *any* [population] which involves only N people, no matter how high their lifetime utilities becomes.” (2018, p. 10, my emphasis)

Pivato is saying that, to avoid Nozick’s Utility Monster, any axiology should respect the following adequacy condition: for any finite population size (no matter how big), there is a population that is better than any population with that size.⁸

The prioritarian theory I will defend in the next section respects the No utility monster axiom and, in some sense, it is a Rank-additive theory. Pivato claims that views such as The Total View and Prioritarianism are Rank-additive theories, because although different transformations can be applied to different entries in this ranking, different transformations do not have to be applied to different entries. The Total View

⁸ While I am sympathetic to views that have this feature, I find it a bit too strong: violating such an axiom would not be a fatal flaw for a theory, indeed this feature might be a cost of a theory. It is crucial for the power of the Monstrous Conclusion as a fundamental constraint on population axiologies that it is limited to cases involving one individual (or perhaps a small sized population). It would not be Monstrous to say that there is some size where, for any population, there is some better population of that size, say involving billions of flourishing lives. In this sense, avoiding the Monstrous Conclusion is a more plausible adequacy condition for theories of population ethics than Pivato’s No utility monsters axiom.

and Prioritarianism are therefore limit cases of rank-additive theories where the rank of a life in the population never makes a difference to the transformations applied to individual's welfare (Pivato 2020 p. 865-866). However, in order to avoid Nozick's Utility Monster, Pivato prefers Rank-additive theories where the rank of a life in the population *does* makes a difference to the transformations function of its wellbeing. In what remains of this section, I show why this should not be considered satisfying.

Pivato claims that, by comparison to the 'trite' observation that Prioritarian views which involve an upper bound can avoid Utility Monsters (2018 p. 10), his Rank-additive theories can have different bounds for different sizes of population and not just a multiplier of the Prioritarian's upper bound. However, this added sophistication can only be achieved by falling subject to a devastating set of further costs. In particular, whenever Rank-additive theories involve a substantive role for an individual's rank, they violate, in especially implausible ways, what is known as Existence Independence. According to Existence Independence:

the ethical evaluation of outcomes concerning some collection K of individuals (say, those currently alive on planet Earth) should not depend upon information about the lifetime utilities —or even the existence —of people outside of K (say, people who died long ago, who will be born in the far future, or who live on other planets) (Pivato 2020, p. 880).

A classic example of an objection targeting violations of Existence Independence was given by Parfit (1986, p. 420). He points out that the value of a life “depends on facts about all previous lives. If the Ancient Egyptians had a very high quality of life, it is more likely to be bad to have a child now. [...] But research in Egyptology cannot be relevant to our decision whether to have children.” It is very hard to believe that the value of a life depends on how the quality of life of people on distant planets, in distant past, or distant future.

In some ways the problem for Rank-additive axiologies is even worse, as Pivato says “Rank-additive axiologies violate Existence Independence in a [...] fundamental way” (2020, p. 880). A feature like the average welfare is not enough to determine an individual's contributory value, on Rank-additive axiologies “we don't even know how to assign ranks to the members of K until we know the lifetime utilities of all the other people not in K ” (Pivato 2020, p. 880). In other words, knowing the average welfare of

the ancient Egyptians would be nowhere near sufficient, we would need to know how many people have and ever will live as well as the complete distribution of all of those people's welfare to even determine an individual's rank, otherwise an individual's contributory value is undefined (Blackorby et al., 2005, sec. 5.1.1; Pivato, 2018, p. 11).⁹

While Pivato expresses willingness to pay this intellectual cost, I agree with many others that the violation of Existence Independence is a sufficient reason to reject a theory, let alone such an extreme violation (Broome, 2004, pp. 117–131; Mulgan, 2001; Nebel, 2022, p. 11; Parfit, 1984, p. 420; Thomas, 2022).

Consequently, I argue instead for a Rank-additive axiology which avoids the Monstrous Conclusion, does not impose an arbitrary threshold, and does not violate Existence Independence. On this theory, an individual's contributory value is bounded above to avoid the Monstrous Conclusions and an individual's rank plays no role in determining the transformation that is applied to their welfare, to avoid violating Existence Independence. I call this theory Asymptotic Prioritarianism.

3. Asymptotic Prioritarianism and the theoretical costs of Prioritarianism

3.1 Prioritarianism

Prioritarianism's central claim is that there is greater moral importance, or priority, to giving welfare increases to the less well off. In other words, there is less moral importance to giving welfare increases to the better off (Parfit 1991; 2012; Holtug 2017; Adler 2021).

While Parfit intended Prioritarianism to be limited to fixed populations (Parfit, 2012, p. 440; Segall, 2022), Prioritarianism, including the kind of Prioritarianism I am about to defend in this section, is compatible with many population axiologies. For now, let us assume its most popular version (Adler, 2019; Holtug, 2017, 2022; Segall, 2022), Total Prioritarianism, according to which population value is the total sum of individual priority-adjusted welfare (we will see other prioritarian population axiologies in section 3.3).

⁹ Pivato (2018) distinguishes between “actualist” and “possibilist” rank-additive theories, each with their own challenges. This distinction does not matter for the present paper, because as he admits in Section 4, Existence Independence fails on both kinds of theories.

To better understand the arguments of this section, some formalization can be useful. Total Prioritarianism can be expressed as ranking outcomes according to the formula $\sum_i g(w_i)$, where w_i is the amount of welfare of each individual i , and $g(\cdot)$ is a concave transformation function. A concave transformation function is a function where the gradient of the slope is always decreasing. A transformation function represents the relation between an individual's welfare and that individual's contributory value to a population. The gradient of the slope of a transformation function represents how much priority we give to an increase of people's welfare when they are better off rather than worse off, that is, the contributory (or marginal) value of each increase in welfare. If the gradient of the slope increases, then the contributory value of additional welfare increases as an individual welfare is higher. If the gradient of the slope *decreases*, as in prioritarian functions, then the contributory value of additional welfare *decreases* as an individual welfare is higher. In other words: with a constantly decreasing value of the slope, the lower an individual's welfare, the more an increase to that individual's welfare contributes to population value.

Prioritarianism is closely related to the Pigou-Dalton principle, the foundation of economic work on inequality (Adler 2013). The Pigou-Dalton principle states that any non-rank-switching fixed transfer from the better off to the less well-off makes a population better. By "non-rank-switching" I mean that "the one who starts out with less does not end up with more than the other" (Adler 2013, p. 1). Any concave prioritarian transformation function will satisfy the Pigou-Dalton principle, as the average gradient for the loss of welfare for the better off individual will be less than the average gradient for the gain of welfare for the less well-off individual.

In addition to fixed transfers, the foundational motivations for Prioritarianism also motivate to transfer welfare from the better off to the worse off even if a fraction of the transferred welfare is lost. This is not a necessary feature of theories respecting Pigou-Dalton, as a theory can use Pigou-Dalton only as a tie-breaker. However, without allowing that a transfer in welfare to the worse off is worth some welfare cost, these theories are to be considered less committed to the core prioritarian intuitions. This intuition prescribes that we should consider it sufficiently more important that the less well-off individual gets the welfare increase that some loss of total welfare can still result in a better population overall. Typically, prioritarians believe that the priority of the worse off is not trivial with respect to wellbeing, it is not just a tie breaker.

For example, if we were to say that it is better to transfer 101 units of welfare from someone with 1,000 units to someone with 10 units, but that it would *not* be better to take 101 units from someone with 1,000 units to give someone starting at 10 an additional 100, then it seems that we would be Prioritarians in name only – there could only be at most trivial amounts of priority between the individuals. We must be willing to pay some specific sufficiently small amount of welfare to make the transfer from the better off person to the worse off person.

This core prioritarian intuition applies so long as one of the individuals is sufficiently better off than the other. Not only it is better to take 101 from someone with 1000 units to give 100 units to someone with 10 units, but also better to take 101 from someone with 100,000 units to give 100 units to someone with 1000 units and so on. To properly respect the core prioritarian intuition, one must satisfy what I call “Non-trivial Priority throughout the Welfare Range”.

Non-trivial Priority throughout the Welfare Range says that there is at least some multiplier of welfare increases such that, for any welfare increase to a less well off individual, there is some better off individual such that it is better to give the initial welfare increase to the less well off individual than the multiplied welfare increase to the better off individual. For example, in the illustrative cases, there was a multiplier of 1.01 such that for a welfare increase (e.g. 100 units) for the less well off individual, it is better that the less well off individual receives the given welfare increase than someone who is 100 times as well off receives the multiplied welfare increase (1.01*100=101 units). I take it to be the case that there is such a multiplier. No matter how well off the better off individual is, in order to fundamentally respect the core Prioritarian intuition, we should prioritise an increase in welfare to the worse off individual over an increase in welfare to the better off individual multiplied by the multiplier.

Now I state *Non-trivial Priority throughout the Welfare Range* more formally. Let w_j and w_i be two welfare levels, m a multiplier greater than 1, Δw a positive change in welfare, and $g(\cdot)$ the prioritarian concave transformation function.

Non-trivial Priority throughout the Welfare Range: $\exists m > 1$ such that $\forall \Delta w$ and $\forall w_i, \exists w_j > w_i$ such that $g(w_i + \Delta w) - g(w_i) > g(w_j + m\Delta w) - g(w_j)$

There is some multiplier of m welfare increase such that for any size of welfare increase Δw and any welfare level w_i , there is some greater welfare level w_j (we can imagine it much greater), where it is better if a less well off individual at welfare w_i receives the welfare increase Δw than if a better off individual at welfare level w_j receives the multiplied welfare increase $m\Delta w$.

While intuitive, *Non-trivial Priority throughout the Welfare Range* does come with an unavoidable cost, as multiple iterations of it imply what I call the Trade-off Condition.

Roughly speaking, where Non-trivial Priority requires that we should be willing to lose some welfare in order to transfer welfare between two sufficiently differently off individuals, the Trade-off Condition states that *any* amount of welfare will be lost in order to transfer welfare between two sufficiently differently off individuals.

More precisely, let Δw be a welfare difference, let w_i and w_j be two wellbeing levels, and let k be a multiplier.

$$\text{Trade-off Condition: } \forall k > 1 \forall \Delta w, \forall w_i, \exists w_j > w_i \text{ such that } g(w_i + \Delta w) - g(w_i) > g(w_j + k\Delta w) - g(w_j)$$

The Trade-off Condition means that, for any multiplier k , any welfare increase Δw and any welfare level w_i , there is some sufficiently better off welfare level w_j such that giving the welfare increase Δw to the person at the given welfare level w_i is better than giving the multiplied welfare increase to the better off individual (even though the multiplier may be arbitrarily large).

To see that *Non-trivial Priority throughout the Welfare Range* implies the Trade-off Condition, consider the following argument.

By *Non-trivial Priority throughout the Welfare Range*, there is some multiplier $m > 1$ such that, for any welfare increase Δw and any initial welfare level w_0 , there is some better off individual at some welfare level w_1 such that it is better to give the given welfare increase Δw to the person at the given welfare level w_0 than the multiplied welfare increase $m\Delta w$ to the better off individual at w_1 .

As *Non-trivial Priority throughout the Welfare Range* applies to any initial welfare level and size of welfare increase, we can apply the principle to welfare level

w_1 and welfare increase $m\Delta w$. By *Non-trivial Priority throughout the Welfare Range*, there is some better off individual w_2 such that it is better to give the given welfare increase $m\Delta w$ to w_1 than the multiplied welfare increase $m^2\Delta w$ to w_2 .

More generally *Non-trivial Priority throughout the Welfare Range* implies that there is a sequence of welfare levels such that for n iterations, it is better to give $m^n\Delta w$ welfare increase to w_n than welfare increase of size $m^{n+1}\Delta w$ to someone at welfare level to someone w_{n+1} . By transitivity, it is better to give a welfare increase of size Δw to someone at w_0 than to give someone a welfare increase of size $m^n\Delta w$ at welfare level w_n . However, as this is true for all n , and because $m > 1$, as n tends to infinity, m^n also tends to infinity. In other words, m^n is unbounded. Therefore, for every $k > 1$, there exists some n such that $m^n > k$. As more welfare is always better, for such an n , giving a welfare increase of size $m^n\Delta w$ is better than giving a welfare increase of size $k\Delta w$.

Finally, we can see that, for any size of $k > 1$, there is some step in the sequence such that, it is better to give a welfare increase of size $m^n\Delta w$ than of size $k\Delta w$ to w_n but better still to give welfare increase of size Δw to welfare level w_0 - an instance of the Trade-off Condition. As Δw and w_0 were chosen arbitrarily, the Trade-off Condition applies to all welfare increases and initial welfare level.

I find the Trade-off Condition troubling – that there is no bound to the amount of welfare that might be sacrificed for the sake of priority is not a feature that would attract me to a theory. Nevertheless, when required to choose between satisfying *Non-trivial Priority throughout the Welfare Range* and avoiding the Trade-off Condition, I accept the Trade-off Condition as a necessary cost of satisfying our core Prioritarian intuition.¹⁰

3.2 Asymptotic Prioritarianism

Having argued that, as Prioritarians, we ought to already be committed to the Trade-off Condition for reasons independent of the Monstrous Conclusion, I now argue that Asymptotic Prioritarianism's strength above and beyond the Trade-off Condition is a small price to pay to avoid the Monstrous Conclusion.

¹⁰ See (Nebel & Stefánsson, 2023) for a discussion of other 'calibration' problems that affect Prioritarian theories.

Asymptotic Prioritarianism is a form of Prioritarianism where the concave transformation function $g(\cdot)$ has an upper bound, or asymptote, corresponding to some population which is better than any one individual population. The contributive value of one individual can always increase the more their welfare increases, but not beyond some limit. This limit corresponds to the value of some population which is better than any one individual population.

Consequently, as the Monster gets better and better off, their contributory value gets closer and closer to the asymptote. Therefore, as the contributory value that the Monster would have with any further welfare increase is bounded above by the asymptote, the moral importance of giving the monster any further welfare increase tends to 0. In other words, Asymptotic Prioritarianism implies the following Absolute Priority Condition. Let Δw be a welfare difference, let w_i and w_j be two wellbeing levels, and let m be a multiplier.

$$\textit{Absolute Priority Condition:} \forall w_i, \forall \Delta w, \exists w_j > w_i \text{ such that } \forall m > 1, g(w_i + \Delta w) - g(w_i) > g(w_j + m\Delta w) - g(w_j)$$

The *Absolute Priority Condition* says that, for any amount of moral importance generated by giving some welfare increase Δw to some individual at welfare level w_i , there is some vastly better off individual at welfare level w_j such that, no matter how big the multiplier m of the welfare increase Δw is, it is more important to give Δw to w_i , than $m\Delta w$ to w_j . In other words, there is a form of absolute priority where it is better to give some welfare increase Δw to a given individual than to give any size of welfare increase $m\Delta w$ to someone who is sufficiently better off (since the better off person is so close to the asymptote, no increase for them could be of more moral importance than the given welfare increase to the given individual).

As with the Trade-off Condition, I take the Absolute Priority Condition to be a counterintuitive implication of an axiology – because of the same central feature that there is no bound to the amount of welfare that might be sacrificed for the sake of priority. The key difference between the Trade-off Condition and the Absolute Priority Condition is just one of quantification order. The Trade-off Condition chooses the arbitrarily large amounts of welfare to be outweighed by priority first, and then ‘finds’ two different welfare levels such that it is more important to give the smaller benefit to

who lives at the worse welfare level than the greater benefit to who lives at the better welfare level. By contrast, the Absolute Priority Condition chooses a benefit to a specific individual, and then ‘finds’ a sufficiently well-off person where there is no amount that could be given to the better-off person which is more morally important than the benefit to the worse-off person.¹¹

While both conditions imply that an arbitrarily large amount of welfare can be outweighed by priority, only the Absolute Priority Condition, which puts a bound to the importance of increasing welfare level, is sufficient to ensure that there is a population which is better than any population only containing the monster – that is, is sufficient to avoid the Monstrous Conclusion. In some ways, without committing to any arbitrary threshold, it captures the sufficientarian intuition that, if we have a significant welfare increase that we can give to some less well-off individual, we do not need to know how much welfare is at stake for some sufficiently better off individual, there is a form of ‘absolute’ priority for this welfare increase to this less well-off individual. When working out whether to give a meal to a hungry child, do we really need to know how much better off Scrooge McDuck could be made with the same amount of resources?

Additionally, something similar to the Absolute Priority Condition is *necessary* to avoid the Monstrous Conclusion. To see why, consider the following argument.

If the Monstrous Conclusion is false, then there is some given population that is not worse than the Monster, no matter how high the Monster’s welfare is. Whatever the features of this given population, there is a possible population that consists in the same individuals of this given population, but is *worse* than the Monster (perhaps they all have lives of suffering).

By increasing one person’s welfare by a small amount Δw , then another, and another, we can gradually transition from the possible population worse than the Monster to the given population which is not worse than the Monster. This means that there is some finite number of small improvements Δw to each individual’s welfare that is the difference between the population worse than some Monster and the population which is not worse than any Monster.

¹¹ I am deeply grateful to an anonymous reviewer for the valuable suggestion to highlight the difference between the Trade-off Condition and the Absolute Priority Condition.

Now, suppose that we have a sequence of comparisons. On the one hand, we have each member of the sequence of small improvements Δw from the larger population that is worse than some Monster to the larger population that is not worse than any Monster. On the other hand, for every step in the sequence, we dramatically increase the Monster's welfare by some amount $m\Delta w$. At the start of the sequence, the larger population is worse than the Monster but by the end of the sequence, the larger population is not worse than the much better off Monster. This is only possible if, for at least some step in the sequence, the small improvement Δw to the larger population has at least as great contributory value than the large welfare increase $m\Delta w$ for the Monster. That is, for some amount of moral importance generated by giving some welfare increase Δw to some individual at welfare level w_i , there is some vastly better off individual at welfare level w_j such that, no matter how big the multiplier m of the welfare increase Δw is, it is at least as important to give Δw to w_i , than $m\Delta w$ to w_j .

The most controversial feature of the Absolute Priority Condition is that fixed welfare increases, no matter how small, for fixed welfare level individuals are more important than any size of welfare increase for some sufficiently well off individual. As the argument above shows, a condition with this order of quantification is required to avoid the Monstrous Conclusion: there must be some step in the argument where arbitrarily small welfare increase, Δw , is more important to give to some w_i than any size of welfare increase $m\Delta w$ to w_j . The Absolute Priority Condition is stronger than the minimum required to avoid the Monstrous Conclusion as the Absolute Priority Condition applies to all given welfare levels. However, this is the only way to avoid some hard to justify threshold between those welfare levels to which it applies and those to which it does not - precisely the kind of issue which made Sufficientarianism implausible.

We have seen that the Absolute Priority Condition is a stronger version of the Trade-off Condition. However, if we are prioritarrians, we already embrace the Trade-off Condition, and thus, are already committed to the idea that there is no bound to the amount of welfare that might be sacrificed for the sake of priority. As prioritarrians, we have two options: either accepting only the Trade-off Condition together with the Monstrous Conclusion, or avoiding the Monstrous Conclusion and accepting the stronger Absolute Priority Condition. I contend that accepting the Absolute Priority Condition rather than the Trade-off Condition is well worth being able to avoid the

Monstrous Conclusion. To the extent to which we find it plausible that there is no bound to the amount of welfare that might be sacrificed for the sake of priority, we should consider Asymptotic Prioritarianism a promising theory for avoiding the Monstrous Conclusion.¹²

However, Prioritarianism is considered vulnerable to a particularly implausible version of the Repugnant Conclusion, according to which, for any population A, there is a better Z population consisting of people whose lives are barely good, even if the total welfare of Z is smaller than the total welfare than A (Adler, 2019; Holtug, 2017, 2022; Segall, 2022). Thus, at this point, some may object that, by accepting Asymptotic Prioritarianism, we are avoiding the Monstrous Conclusion by implying a worse version of the Repugnant Conclusion. In the next section, I will show that this is not the case: Asymptotic Prioritarianism can be implemented on different aggregative theories, and some implementations of Asymptotic Prioritarianism can avoid both the Monstrous and the Repugnant Conclusion.

3.3 Different Asymptotic Prioritarianism axiologies

Prioritarianism is a view about the importance of individual wellbeing, and by itself it is not a full population axiology. Parfit considers Prioritarianism a theory for fixed populations, and suggests to ‘quarantine’ it for variable populations precisely because of its implications when combined with other theories in population axiology – notably the Total View (Parfit, 2012, p. 440; Segall, 2022). However, Prioritarianism, and Asymptotic Prioritarianism in particular, can be combined with a wide range of population axiologies. While some combinations are less promising, as Asymptotic Prioritarianism exacerbates existing objections to those population axiologies, other combinations allow for the benefits of both components of the resulting population axiology and, finally, in some cases, Asymptotic Prioritarianism can also help mitigate the existing objections to those views.

¹² There are also Non-trivial Priority theories that avoid the Monstrous Conclusion by having a prioritarian function whose slope suddenly becomes flat, becoming a sharp threshold for contributory value. These theories avoid the Monstrous Conclusion without implying the Absolute Priority Condition. The advantage that Asymptotic Prioritarianism has over these theories, and over any theory that avoids the Monster with a threshold, is identical to the advantage that it has over Sufficieritarian theories.

The most discussed Prioritarian axiology is Total Prioritarianism, where first each individual's welfare is priority-adjusted to give each the moral importance, or more precisely contributory value, of each individual's welfare, and then populations are ranked according to the total of these amounts of contributory value (Adler, 2019; Holtug, 2017, 2022; Segall, 2022). However, while Total Asymptotic Prioritarianism can fare better than other forms of Total Prioritarianism, as it can avoid the Monstrous Conclusion, Total Asymptotic Prioritarianism shares the problems of other forms of Total Prioritarianisms in that it exacerbates an existing objection to the Total View.

To see how Total Asymptotic Prioritarianism can avoid the Monstrous Conclusion is relatively straightforward. First, Asymptotic Prioritarianism implies that there is some limit of contributory value that the Monster cannot exceed. Given that lives worth living have some positive individual contributory value, Total Asymptotic Prioritarianism implies that a sufficient number of those lives will have greater total contributory value than the limit the Monster cannot exceed and, therefore, have greater total contributory value than any Monster. Such a population would be better than any one-individual population, thus avoiding the Monstrous Conclusion.

However, Total Asymptotic Prioritarianism has an even worse form of the classic objection to the Total View, namely the Repugnant Conclusion. The Repugnant Conclusion states that, for any population, there is some better population consisting only of lives that are barely worth living. Total Asymptotic Prioritarianism implies an especially bad form of the Repugnant Conclusion. For any population, that population has some amount of total contributory value. Barely worth living lives will have some low amount of contributory value but sufficiently many of them will have greater total contributory value than any given level. The difference from the classical Total View is that, because of the greater contributory value of the welfare for the less well off than the better off, even fewer barely worth living lives are required to have greater total contributory value than is the case with the classical Total View.

Given its serious problems in dealing with the Repugnant Conclusion, Asymptotic Total Prioritarianism cannot be considered a promising theory for population ethics, even though it avoids the Monstrous Conclusion. Fortunately, Asymptotic Prioritarianism can be combined with population axiological theories that avoid the Repugnant Conclusion.

The most discussed population axiology that avoids the Repugnant Conclusion is the Average View. However, a combination of the Average View and Asymptotic

Prioritarianism is not a promising candidate. While there might be some benefits to combining Asymptotic Prioritarianism with the Average View, such an Average Asymptotic Prioritarianism would be unable to avoid the Monstrous Conclusion, as the welfare of any given population is lower than the average welfare of the Monster alone.

A far more promising candidate is the combination of Asymptotic Prioritarianism with Critical Level and Critical Range Views, a set of views with considerable existing support level (Blackorby et al., 1997, 1998, 2005; Gustafsson, 2020; Rabinowicz, 2009, 2022; Thornley, 2022). Critical Level and Critical Range Asymptotic Prioritarianism take into account both prioritarian considerations for the welfare of the less well off compared to the better off, as well as requiring that welfare levels must be above a critical level or range in order to have a positive contribution to population value. Critical Level and Critical Range Asymptotic Prioritarianism rank populations according to their total welfare, adjusted both by critical level (or range) and asymptotic prioritarian factors.

Critical Level and Critical Range Asymptotic Prioritarianism can avoid the Monstrous Conclusion because there is some limit to the contributory value of one individual's welfare. Given that the welfare of individuals above the critical level or range has positive contributory value, sufficiently many of them will have greater total contributory value than the limit of the contributory value of one single individual's welfare. Hence, that population will be better than any Monster, according to Critical Level and Critical Range Asymptotic Prioritarianism.

Moreover, Critical Level and Critical Range Asymptotic Prioritarianism can avoid the Repugnant Conclusion. For any population with all individuals above the critical level or range, that population will have positive total contributory value. By contrast, any population consisting only of individuals with lives that are barely worth living will be below the critical level or range (or within the range) and therefore not have positive total contributory value. Hence, these populations are not worse than any population consisting only of individuals with lives that are barely worth living – thus avoiding the Repugnant Conclusion.

Critical Level and Critical Range Asymptotic Prioritarianism are therefore able to combine the most important benefits of both Asymptotic Prioritarianism and Critical Level or Critical Range Views. This does not mean that we should necessarily endorse Critical Level or Critical Range Asymptotic Prioritarianism in population ethics, as they combine their problems, too. Critical Level or Critical Range Views have problems

with low welfare level lives, which remain in Critical Level Asymptotic Prioritarianism (for the problems of Critical Level theories, see (Arrhenius, 2000, p. 73; Williamson, 2021). For the problems of Critical Range Theories, see (Broome, 2004, p. 148-170, 2009)). In addition to these problems, Critical Level and Critical Range Asymptotic Prioritarianism need to deal with the Trade-off Condition raised in the previous section.

Another promising candidate is the combination of Asymptotic Prioritarianism with Combined Views. Combined Asymptotic Prioritarianism combines asymptotic diminishing contributive (or marginal) value for both additional welfare and additional individuals at a given welfare. Combined Asymptotic Prioritarianism is thus able to avoid both the Monstrous Conclusion and Repugnant Conclusion. To see this, consider the following.

No matter how many individuals are added with barely worth living lives, the contributory value of the population cannot exceed the limit for barely worth living lives. As the limit is greater for higher amounts of welfare, there is a population containing individuals with the higher amount of welfare which is better than all populations consisting of barely worth living lives. Similarly, no matter how much one individual's welfare is increased, the contributory value of the population cannot exceed the limit for one individual. As the limit for more individuals is higher, there is a population with more people which is better than any population consisting of only one individual. Therefore, Combined Asymptotic Prioritarianism avoids both the Monstrous Conclusion and Repugnant Conclusion.

Moreover, unifying Combined Views with Asymptotic Prioritarianism may mitigate the existing objections against Combined Views. For example, Sider himself is skeptical of his own Combined View since "it generates rather extreme results with respect to distributive justice" (Sider 1991, note 18).¹³ However, when combined with Asymptotic Prioritarianism, there would be a significant corrective to these distributional concerns.

To conclude, I do not endorse any one population axiology, each has its advantages and well-known problems. However, the potential of Asymptotic Prioritarianism in population ethics is enormous: further research is required in population ethics on Trade-off Conditions in Prioritarianism and how Asymptotic

¹³ Note that Sider's view differs from other Combined Views: see footnote 5.

Prioritarianism can combine with other theories in population axiology to mitigate existing problems.

4. Conclusion

In this paper, I introduced the Monstrous Conclusion, a disturbing implication of most theories of population ethics – that for any population there is a better population consisting of one individual. The Monstrous Conclusion is roughly as disturbing as the Repugnant Conclusion, but has received far less attention, and examining it may clarify our intuitions concerning Nozick’s more famous Utility Monster. The Monstrous Conclusion is implied by most prominent theories of population ethics (total and average theories, theories combining total and average aspects, person affecting theories).

I argued that there is one hitherto undefended form of Prioritarianism that is particularly well equipped to avoid the Monstrous Conclusion. This form of Prioritarianism is a member of what I call the family of Non-Trivial Priority theories, according to which it is better to transfer welfare from the better off to the worse off even if some portion of that welfare is lost. Particularly, I show that, on Non-Trivial Priority theories, this lost portion of welfare has to get bigger as the two welfare levels are distant: this implies what I call the Trade-off Condition, according to which, for any two welfare improvements, one greater than the other, there are two welfare levels such that it is better to give the smallest benefit to the person at the lower welfare level than the greater benefit to the person at the greater welfare level. However, those who take seriously the priority view have to agree that the Trade-off Condition is a cost we should be willing to pay.

The form of Non-Trivial Priority that best avoids the Monstrous Conclusion is what I call Asymptotic Prioritarianism, according to which the priority of the welfare of well off people over less well off people cannot exceed some limit. This limits how much priority a single life can get in a population, but does not limit how much priority multiple people can get in a population: this enables Asymptotic Prioritarianism to avoid the Monstrous Conclusion.

Avoiding the Monstrous Conclusion, too, comes at an intellectual cost: I call it the Absolute Priority Principle, according to which, for any welfare level and any welfare increase, there is a greater welfare level such that it is better to give the given welfare level to the smaller welfare level than any welfare increase greater than the

given welfare increase. However, Asymptotic Prioritarianism pays this cost most acceptably, as the Absolute Priority Principle is just a more pronounced version of the Trade-off Condition already implied by Non-Trivial Priority theories, and is not much stronger than a necessary condition to avoid the Monstrous Conclusion. However, by itself Asymptotic Prioritarianism is not a complete population axiology: it must be combined with some further form of aggregation.

I show that combining Asymptotic Prioritarianism with other population axiologies can enable us to combine the positive features of each, although it may not mitigate or even exacerbate their other flaws. For example, although Total Asymptotic Prioritarianism can avoid the Monstrous Conclusion, it faces an even worse form of the Repugnant Conclusion. By contrast, Critical Level and Critical Range Asymptotic Prioritarianism are able to avoid both the Monstrous Conclusion and the Repugnant Conclusion.

Crucially, I do not argue that we should endorse one population axiology as a result of the arguments in this paper. I have deep reservations about whether the Trade-off Conditions facing Non-Trivial Priority theories can be adequately addressed or explained away and existing population axiologies all have well-known counterintuitive consequences. However, avoiding the Monstrous Conclusion is a problem that population axiologies must take seriously and the potential of Asymptotic Prioritarianism and other ways of avoiding the Monstrous Conclusion in population ethics is enormous.

Bibliography

- Adler, M. D. (2019). *Measuring Social Welfare: An Introduction* (1st ed.). Oxford University Press. <https://doi.org/10.1093/oso/9780190643027.001.0001>
- Adler, M. D., & Norheim, O. F. (Eds.). (2022). *Prioritarianism in practice*. Cambridge University Press.
- Arneson, R. J. (2000). Perfectionism and Politics. *Ethics*, 111(1), 37–63. <https://doi.org/10.1086/233418>
- Arneson, R. J. (2002). Why Justice Requires Transfers to Offset Income and Wealth Inequalities. *Social Philosophy and Policy*, 19(1), 172–200. <https://doi.org/10.1017/S0265052502191084>
- Arrhenius, G. (2000). *Future Generations: A challenge for moral theories*. Uppsala University Press.
- Arrhenius, G. (Forthcoming). *Population Ethics: The challenge of future generations*. Oxford University Press.

- Arrow, K. J. (1971). *Essays in the theory of risk-bearing*. Markham Publishing Company.
- Beard, S. (2020). Perfectionism and the Repugnant Conclusion. *The Journal of Value Inquiry*, 54(1), 119–140. <https://doi.org/10.1007/s10790-019-09687-4>
- Beckstead, N., & Thomas, T. (2023). A paradox for tiny probabilities and enormous values. *Noûs*, nous.12462. <https://doi.org/10.1111/nous.12462>
- Beglin, D. (2017). Should I choose to never die? Williams, boredom, and the significance of mortality. *Philosophical Studies*, 174(8), 2009–2028. <https://doi.org/10.1007/s11098-016-0785-1>
- Blackorby, C., Bossert, W., & Donaldson, D. (1997). Critical-Level Utilitarianism and the Population-Ethics Dilemma. *Economics and Philosophy*, 13(2), 197–230. <https://doi.org/10.1017/S026626710000448X>
- Blackorby, C., Bossert, W., & Donaldson, D. (2005). *Population issues in social-choice theory, welfare economics and ethics*. Cambridge University Press.
- Blackorby, C., Bossert, W., Donaldson, D., & Fleurbaey, M. (1998). Critical Levels and the (Reverse) Repugnant Conclusion. *Journal of Economics*, 67(1), 1–15. <https://doi.org/10.1007/BF01227760>.
- Bossert, W., Cato, S., & Kamaga, K. (2022). Critical-level Sufficiencyarianism*. *Journal of Political Philosophy*, 30(4), 434–461. <https://doi.org/10.1111/jopp.12267>
- Briggs, R., & Nolan, D. (2015). Utility Monsters for the Fission Age. *Pacific Philosophical Quarterly*, 96(3), 392–407. <https://doi.org/10.1111/papq.12079>
- Broome, J. (2004). *Weighing Lives* (1st ed.). Oxford University Press Oxford. <https://doi.org/10.1093/019924376X.001.0001>
- Broome, J. (2009). Reply to Rabinowicz. *Philosophical Issues*, 19, 412–417. [10.1111/j.1533-6077.2009.00175.x](https://doi.org/10.1111/j.1533-6077.2009.00175.x)
- Brown, C. (2005). Priority or Sufficiency ...or both? *Economics and Philosophy*, 21(2), 199–220. <https://doi.org/10.1017/S0266267105000568>
- Bruckner, D. W. (2012). Against the Tedium of Immortality. *International Journal of Philosophical Studies*, 20(5), 623–644. <https://doi.org/10.1080/09672559.2012.713383>
- Buchak, L. (2017). *Risk and rationality* (First paperback edition). Oxford University Press.
- Bykvist, K. (2007). The Benefits of Coming into Existence. *Philosophical Studies*, 135(3), 335–362. <https://doi.org/10.1007/s11098-005-3982-x>
- Carlson, E. (1998). Mere Addition and Two Trilemmas of Population Ethics. *Economics and Philosophy*, 14(2), 283–306. <https://doi.org/10.1017/S0266267100003862>
- Casal, P. (2007). Why Sufficiency Is Not Enough. *Ethics*, 117(2), 296–326. <https://doi.org/10.1086/510692>
- Cowen, T., & High, J. (1988). Time, bounded utility, and the St. Petersburg paradox. *Theory and Decision*, 25(3), 219–223. <https://doi.org/10.1007/BF00133163>
- Crisp, R. (2003). Equality, Priority, and Compassion. *Ethics*, 113(4), 745–763. <https://doi.org/10.1086/373954>
- Dillenberger, D., & Vijay Krishna, R. (2014). Expected utility without bounds—A simple proof. *Journal of Mathematical Economics*, 52, 143–147. <https://doi.org/10.1016/j.jmateco.2013.12.011>
- Dorsey, D. (2014). Equality-tempered Prioritarianism. *Politics, Philosophy & Economics*, 13(1), 45–61. <https://doi.org/10.1177/1470594X13483479>
- Fischer, J. M., & Mitchell-Yellin, B. (2014). Immortality and Boredom. *The Journal of Ethics*, 18(4), 353–372. <https://doi.org/10.1007/s10892-014-9172-3>
- Fishburn, P. C. (1976). Unbounded Utility Functions in Expected Utility Theory. *The*

Quarterly Journal of Economics, 90(1), 163. <https://doi.org/10.2307/1886093>

Frankfurt, H. (1987). Equality as a Moral Ideal. *Ethics*, 98(1), 21–43. <https://doi.org/10.4324/9781315199795-15>

Gorman, A. G. (2017). Williams and the Desirability of Body-Bound Immortality Revisited. *European Journal of Philosophy*, 25(4), 1062–1083. <https://doi.org/10.1111/ejop.12184>

Greaves, H. (2017). Population axiology. *Philosophy Compass*, 12(11), e12442. <https://doi.org/10.1111/phc3.12442>

Greene, P. (2017). Value in Very Long Lives. *Journal of Moral Philosophy*, 14(4), 416–434. <https://doi.org/10.1163/17455243-46810057>

Grill, K. (2023). The Sum of Averages: An Egyptology-Proof Average View. *Utilitas*, 35(2), 103–118. <https://doi.org/10.1017/S0953820822000450>

Gustafsson, J. E. (2020). Population axiology and the possibility of a fourth category of absolute value. *Economics and Philosophy*, 36(1), 81–110. <https://doi.org/10.1017/S0266267119000087>

Gustafsson, J. E. (2022). Our Intuitive Grasp of the Repugnant Conclusion. In G. Arrhenius, K. Bykvist, T. Campbell, & E. Finneron-Burns (Eds.), *The Oxford Handbook of Population Ethics*, 370–389. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190907686.013.25>

Hájek, A., & Smithson, M. (2012). Rationality and indeterminate probabilities. *Synthese*, 187(1), 33–48. <https://doi.org/10.1007/s11229-011-0033-3>

Heyd, D. (1988). Procreation and value can ethics deal with futurity problems? *Philosophia*, 18(2–3), 151–170. <https://doi.org/10.1007/BF02380074>

Hirose, I. (2016). Axiological Sufficiency. In C. Fourie & A. Rid (Eds.), *What is Enough?*, 51–68. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199385263.003.0004>

Holtug, N. (2010). *Persons, interests, and justice*. Oxford University Press.

Holtug, N. (2017). Prioritarianism. In N. Holtug, *Oxford Research Encyclopedia of Politics*. Oxford University Press. <https://doi.org/10.1093/acrefore/9780190228637.013.232>

Holtug, N. (2022). Prioritarianism and Population Ethics. In G. Arrhenius, K. Bykvist, T. Campbell, & E. Finneron-Burns (Eds.), *The Oxford Handbook of Population Ethics*, 37–62. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190907686.013.12>

Huemer, M. (2008). In Defence of Repugnance. *Mind*, 117(468), 899–933. [10.1093/mind/fzn079](https://doi.org/10.1093/mind/fzn079)

Hurka, T. (1983). Value and Population Size. *Ethics*, 93(3), 496–507. <https://doi.org/10.1086/292462>

Kagan, S. (2012). *Death*. Yale University Press.

Kamm, F. (2015). Cost Effectiveness Analysis and Fairness. *Journal of Practical Ethics*, 3(1), 1–14.

Kitcher, P. (2000). Parfit’s Puzzle. *Noûs*, 34(4), 550–577. <https://doi.org/10.1111/0029-4624.00278>

Kosonen, P. (2022). *Tiny Probabilities of Vast Value* (PhD dissertation at Oxford University).

Kreps, D. M. (2013). *Microeconomic foundations. 1: Choice and competitive markets*. Princeton University Press

McTaggart, J. M. E., & Broad, C. D. (1968). *The nature of existence. Vol. 2* (Reprinted). At the University Press.

Miller, D. (2021). Justice. In *The Stanford Encyclopedia of Philosophy* (Fall 2021

Edition). Edward N. Zalta. <https://plato.stanford.edu/archives/fall2021/entries/justice/>

Mogensen, A. L. (2022). Against Large Number Scepticism. In Jeff McMahan, and others (eds) *Ethics and Existence: the Legacy of Derek Parfit*, 311–330. Oxford University Press. <https://doi.org/10.1093/oso/9780192894250.003.0012>

Mulgan, T. (2001). What’s Really Wrong with the Limited Quantity View? *Ratio*, 14(2), 153–164. <https://doi.org/10.1111/1467-9329.00152>

Nebel, J. M. (2017). Priority, Not Equality, for Possible People. *Ethics*, 127(4), 896–911. <https://doi.org/10.1086/691568>

Nebel, J. M. (2022). Totalism without Repugnance. In In Jeff McMahan, and others (eds) *Ethics and Existence: the Legacy of Derek Parfit*, 200–231. Oxford University Press. <https://doi.org/10.1093/oso/9780192894250.003.0009>

Nebel, J. M., & Stefánsson, H. O. (2023). Calibration dilemmas in the ethics of distribution. *Economics and Philosophy*, 39(1), 67–98. <https://doi.org/10.1017/S0266267121000298>

Ng, Y.-K. (1989). What Should We Do About Future Generations?: Impossibility of Parfit’s Theory X. *Economics and Philosophy*, 5(2), 235–253. <https://doi.org/10.1017/S0266267100002406>

Nover, H., & Hájek, A. (2004). Vexing Expectations. *Mind*, 113(450), 237–249. <https://doi.org/10.1093/mind/113.450.237>

Nozick, R. (1974). Anarchy, state, and utopia. Basic Books

Otsuka, M. (2012). Prioritarianism and the Separateness of Persons. *Utilitas*, 24(3), 365–380. <https://doi.org/10.1017/S0953820812000064>

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

Parfit, D. (1991). *Equality or Priority*. University of Kansas, Department of Philosophy.

Parfit, D. (2004). Overpopulation and the Quality of Life. In T. Tännsjö & J. Ryberg (Eds.), *The Repugnant Conclusion*, 7–22. Springer Netherlands. https://doi.org/10.1007/978-1-4020-2473-3_2

Parfit, D. (2012). Another Defence of the Priority View. *Utilitas*, 24(3), 399–440. <https://doi.org/10.1017/S095382081200009X>

Parfit, D. (2016). Can We Avoid the Repugnant Conclusion? *Theoria*, 82(2), 110–127. <https://doi.org/10.1111/theo.12097>

Pivato, M. (2014). Additive representation of separable preferences over infinite products. *Theory and Decision*, 77(1), 31–83. <https://doi.org/10.1007/s11238-013-9391-2>

Pivato, M. (2020). Rank-Additive Population Axiology. *Economic Theory* 69, 861–918. <https://doi.org/10.1007/s00199-019-01194-8>

Pressman, M. (2015). A Defence of Average Utilitarianism. *Utilitas*, 27(4), 389–424. <https://doi.org/10.1017/S0953820815000072>

Pummer, T. (2013). Intuitions about large number cases. *Analysis*, 73(1), 37–46. <https://doi.org/10.1093/analysis/ans134>

Rabinowicz, W. (2009). Broome and the Intuition of Neutrality. *Philosophical Issues*, 19, 389–411. <https://doi.org/10.1111/j.1533-6077.2009.00174.x>

Rabinowicz, W. (2022). Getting Personal: The Intuition of Neutrality Reinterpreted. In G. Arrhenius, K. Bykvist, T. Campbell, & E. Finneron-Burns (Eds.), *The Oxford Handbook of Population Ethics*, 113–141. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190907686.013.17>

Roberts, M. A. (2009). The nonidentity problem and the two envelope problem: When is one act better for a person than another? In M. A. Roberts & D. T. Wasserman (Eds.), *Harming future persons: Ethics, genetics and the nonidentity*

- problem*, 201-228. Springer.
- Rosenqvist, S. (2020). *Hedonistic Act Utilitarianism. Action Guidance and Moral Intuitions*. Uppsala University, Department of Philosophy.
- Ross, J. (2015). Rethinking the Person-Affecting Principle. *Journal of Moral Philosophy*, 12(4), 428–461. <https://doi.org/10.1163/17455243-01204004>
- Savage, L. J. (1954). *The foundations of statistics*. Wiley
- Segall, S. (2022). Quarantined Prioritarianism. In Jeff McMahan, and others (eds) *Ethics and Existence: the Legacy of Derek Parfit*, 552–576. Oxford University Press. <https://doi.org/10.1093/oso/9780192894250.003.0021>
- Sider, T. R. (1991). Might theory X be a theory of diminishing marginal value? *Analysis*, 51(4), 265–271. <https://doi.org/10.1093/analys/51.4.265>
- Smuts, A. (2011). Immortality and Significance. *Philosophy and Literature*, 35(1), 134–149. <https://doi.org/10.1353/phl.2011.0002>
- Tännsjö, T. (2002). Why We Ought to Accept the Repugnant Conclusion. *Utilitas*, 14(3), 339–359. <https://doi.org/10.1017/S0953820800003642>
- Temkin, L. S. (2008). Is Living Longer Living Better? *Journal of Applied Philosophy*, 25(3), 193–210. <https://doi.org/10.1111/j.1468-5930.2008.00411.x>
- Temkin, L. S. (2012). *Rethinking the good: Moral ideals and the nature of practical reasoning*. Oxford university press.
- Thomas, T. (2018). Some Possibilities in Population Axiology. *Mind*, 127(507), 807–832. <https://doi.org/10.1093/mind/fzx047>
- Thomas, T. (2022). Separability and Population Ethics. In G. Arrhenius, K. Bykvist, T. Campbell, & E. Finneron-Burns (Eds.), *The Oxford Handbook of Population Ethics*, 270–295. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190907686.013.21>
- Thornley, E. (2022). Critical Levels, Critical Ranges, and Imprecise Exchange Rates in Population Axiology. *Journal of Ethics and Social Philosophy*, 22(3). <https://doi.org/10.26556/jesp.v22i3.1593>
- Timmer, D. (2022). Justice, Thresholds, and the Three Claims of Sufficiencyarianism*. *Journal of Political Philosophy*, 30(3), 298–323. <https://doi.org/10.1111/jopp.12258>
- Vallentyne, P. (1991). The Problem of Unauthorized Welfare. *Noûs*, 25(3), 295. <https://doi.org/10.2307/2215505>
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.
- Voorhoeve, A., & Fleurbaey, M. (2016). Priority or Equality for Possible People? *Ethics*, 126(4), 929–954. <https://doi.org/10.1086/686000>
- Williams, B. (1973). *Problems of the Self: Philosophical Papers 1956–1972* (1st ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511621253>
- Williamson, P. (2021). A New Argument Against Critical-Level Utilitarianism. *Utilitas*, 33(4), 399–416. <https://doi.org/10.1017/S0953820821000133>