

Critical Levels, Critical Ranges, and Imprecise Exchange Rates in Population Axiology

*Elliott Thornley**

Forthcoming in the *Journal of Ethics and Social Philosophy*

Abstract: According to Critical-Level Views in population axiology, an extra life improves a population only if that life's welfare exceeds some fixed 'critical level.' An extra life at the critical level leaves the new population equally good as the original. According to Critical-Range Views, an extra life improves a population only if that life's welfare exceeds some fixed 'critical range.' An extra life within the critical range leaves the new population incommensurable with the original.

In this paper, I sharpen some old objections to these views and offer some new ones. Critical-Level Views cannot avoid certain Repugnant and Sadistic Conclusions. Critical-Range Views imply that lives featuring no good or bad components whatsoever can nevertheless swallow up and neutralise goodness or badness. Both classes of view entail that certain small changes in welfare correspond to worryingly large differences in contributive value.

I then offer a view that retains much of the appeal of Critical-Level and Critical-Range Views while avoiding the above pitfalls. On the Imprecise Exchange Rates View, the quantity of some good required to outweigh a given unit of some bad is imprecise. This imprecision is the source of incommensurability between lives and populations.

1. Introduction

How do we determine whether one population is at least as good as another? Here is one easy answer. We use a number to represent each person's welfare – how good their life is for them – with the size of the number proportional to how good their life is. Positive numbers represent good lives, negative numbers represent bad lives, and zero represents lives that are neither good nor bad. We then sum these numbers to get the value of each population. A population X is at least as good as a population Y iff the value of X is at least as great as the

* Comments and questions welcome at elliott.thornley@philosophy.ox.ac.uk

value of Y . A theory of how populations relate with respect to goodness is called a *population axiology*, and we can call this population axiology the *Total View*.

The Total View implies that we can improve populations by adding lives that are barely worth living, and some find this implication distasteful. We can avoid this implication by first subtracting some positive constant from the number representing a person's welfare and then summing the results. Call these population axiologies *Critical-Level Views*.

Critical-Level Views cannot account for two intuitions that many people find appealing. The first is that there is a *range* of welfare levels such that adding lives at these levels makes a population neither better nor worse. The second is that populations of different sizes may be *incommensurable*, so that neither population is better than the other and yet nor are they equally good. In that case, we might prefer to calculate the value of each population on a range of critical levels. We can then claim that X is at least as good as Y iff the value of X is at least as great as the value of Y relative to each level within this range. If neither X nor Y is at least as good as the other, they are incommensurable. Call these population axiologies *Critical-Range Views*.

Critical-Level and Critical-Range Views fall within the more general class of *Critical-Set Views*. I offer a characterisation and taxonomy of these views below, along with six objections that tell against various views in this taxonomy. Some views imply Repugnant or Sadistic Conclusions. Other views make neutrality implausibly greedy. Each view implies that certain small differences in welfare correspond to worryingly large differences in how a life affects the value of a population, and no view can account for the incommensurability between lives and same-size populations without extra theoretical resources.

I then offer a view that retains much of the appeal of Critical-Set Views while avoiding many of the above pitfalls. The *Imprecise Exchange Rates View* has its start in the observation that there are often no precise truths about whether it is worth undergoing some bad for the sake of some good. It makes sense of this observation by claiming that various *exchange rates* between goods and bads are imprecise. This imprecision renders certain combinations of goods and bads incommensurable with other combinations. The view thus provides a natural explanation of incommensurability between lives and same-size populations, avoids the Sadistic Conclusion along with the most concerning instances of Repugnance and Greediness, and has many other advantages besides.

I characterise and taxonomise Critical-Set Views in Section 2, and object to them in Section 3. I introduce the Imprecise Exchange Rates View in Section 4, canvas its advantages in Section 5, and address some objections in Section 6. I sum up in Section 7.

2. Critical-Set Views

Foundational to Critical-Set Views is the notion of a *life*. I follow Broome (2004, 94–95) in loosely defining a life as ‘how things are for a person,’ where this phrase is understood to include all those things that can affect that life’s *welfare*, how good the life is for the person living it. This definition jars somewhat with our ordinary understanding of a life. Depending on our theory of welfare, it might count events occurring after a person’s death as part of their life. But for our purposes, this terminological strangeness is of little consequence. The definition also allows that more than one person can live the same life. This possibility simplifies the ensuing discussion.

Advocates of Critical-Set Views assume that welfare is *measurable on an interval scale* and *interpersonally comparable*. Measurability on an interval scale allows us to talk meaningfully about ratios of differences in welfare, so that claims like the following are meaningful: ‘The difference in welfare between Ada’s life as an artist and Ada’s life as a baker is twice the size of the difference in welfare between Ada’s life as a baker and Ada’s life as a consultant.’ Interpersonal comparability allows us to compare the welfare of different people, so that claims like the following are meaningful: ‘Ada’s life as an artist contains more welfare than Bob’s life as a baker.’ This claim is equivalent to the claim that ‘Ada’s life as an artist is *personally better* than Bob’s life as a baker.’ In other words, ‘Ada’s life as an artist is better *for her* than Bob’s life as a baker is *for him*.’ I mostly use the terminology of personal betterness below.

Most advocates of Critical-Set Views claim that each life’s welfare can be represented by a real-valued function w , so that a life x is at least as personally good as a life y iff $w(x) \geq w(y)$, and the difference in welfare between x and y is k times the difference in welfare between y and z iff $|w(x) - w(y)| = k|w(y) - w(z)|$.¹ This assumption implies that each pair of lives is *commensurable* with respect to welfare. That is, for all possible lives x and y , x is at least as personally good as y or y is at least as personally good as x .

Critical-Set Views typically go on to sort lives into absolute categories. Which category a life falls in depends on how it compares to some standard: a life is *personally good* iff it is better than the standard, *personally bad* iff it is worse than the standard, and *personally neutral* iff it is neither better nor worse than the standard. The category of personally neutral lives can be refined further: a life is personally *strictly* neutral iff it is equally good as the standard and personally *weakly* neutral iff it is incommensurable with the standard. The standard in question is defined differently by different authors. Some define it as nonexistence (Arrhenius and Rabinowicz 2015a). Others define it as a life

¹ Gustafsson (2020) is one exception. I discuss his view below.

constantly at a neutral level of temporal welfare (Broome 2004, 68; Bykvist 2007, 101). Still others define it as a life without any good or bad components: features of a life that are good or bad for the person living it (Arrhenius 2000b, 26). With one caveat, Critical-Set Views are compatible with each definition.²

So much for comparing lives. Comparing populations requires more assumptions. The first thing to note is that populations are compared with respect to *general* goodness, rather than personal goodness. Broome's *Principle of Personal Good* (2004, 120) links the two:

The Principle of Personal Good

Take two populations X and Y that feature all the same people. If each person's life in X is equally personally good as their life in Y , then X and Y are equally generally good. If each person's life in X is at least as personally good as their life in Y , and some person's life in X is personally better than their life in Y , then X is generally better than Y .³

This principle implies that only facts about the identities of existing people and their welfare levels are evaluatively significant. Factors like the beauty of the world matter only derivatively if they matter at all. The Principle of Personal Good simplifies the presentation of Critical-Set Views but is not strictly necessary. Advocates can instead append an 'all else equal' clause to their other claims, specifying that all facts besides identities and welfare levels are to be held fixed.

A second assumption of Critical-Set Views is *Anonymity*:

Anonymity

Any two populations that feature the same number of lives at each welfare level are equally good.

This principle implies that identities are evaluatively insignificant. Every person's welfare matters equally, so a population cannot be made better or worse by changing the identity associated with a given welfare level. Anonymity allows us to represent each population with a distribution – a finite, unordered list of welfare levels allowing repetitions – so that one population is at least as good as another iff its distribution is at least as good. Henceforth, $\llbracket X \rrbracket$ denotes the distribution corresponding to population X . Distributions and welfare levels can

² The caveat is that *Neutral-Range Views* – explained below – cannot be paired with the latter two definitions. Neutral-Range Views claim both that all lives are personally commensurable with each other, and that some lives are personally incommensurable with the standard. That means that the standard cannot be a life. I thank an anonymous reviewer for pointing this out.

³ When context makes it obvious that I am comparing populations, I drop the 'generally' from such comparisons.

be concatenated, so that $\llbracket X \rrbracket \cup \llbracket Y \rrbracket$ denotes the distribution composed of all the welfare levels in $\llbracket X \rrbracket$ and $\llbracket Y \rrbracket$ and $\llbracket X \rrbracket \cup w(x)$ denotes the distribution composed of all the welfare levels in $\llbracket X \rrbracket$ plus the welfare level of life x .

This notation is useful in clarifying the notion of a life's *contributive value* relative to a population: how a life contributes to the general goodness of a population. A life x is contributively good (bad/strictly neutral/weakly neutral) relative to a population X iff $\llbracket X \rrbracket \cup w(x)$ is better than (worse than/equally good as/incommensurable with) $\llbracket X \rrbracket$.⁴ To these absolute classifications of contributive value, we can add comparative ones. A life x is contributively better than (worse than/equally good as/incommensurable with) a life y relative to a population X iff $\llbracket X \rrbracket \cup w(x)$ is better than (worse than/equally good as/incommensurable with) $\llbracket X \rrbracket \cup w(y)$.

A third assumption of Critical-Set Views is *Transitivity*:

Transitivity

For any populations X , Y , and Z , if X is at least as good as Y , and Y is at least as good as Z , then X is at least as good as Z .

A fourth is *Critical Set*:

Critical Set

For at least one possible population, there is a set of welfare levels such that lives at these welfare levels are contributively neutral relative to that population. Lives at welfare levels above (below) the set are contributively good (bad) relative to that population.

And a fifth is *Separability over Lives*:

Separability over Lives

For any populations X , Y , and Z , X is at least as good as Y iff $\llbracket X \rrbracket \cup \llbracket Z \rrbracket$ is at least as good as $\llbracket Y \rrbracket \cup \llbracket Z \rrbracket$.

This principle implies that the critical set is the same for all populations.⁵ As a result, each life has the same contributive value relative to all populations. If a life x is contributively good (bad/strictly neutral/weakly neutral) relative to some

⁴ Gustafsson (2020) calls weakly neutral lives ‘undistinguished.’

⁵ Here is how. Let $w(x)$ be some welfare level above the critical set in some population X . By Critical Set, this is so iff $\llbracket X \rrbracket \cup w(x)$ is better than $\llbracket X \rrbracket$. By Separability, this is true iff, for any population Y , $\llbracket X \rrbracket \cup w(x) \cup \llbracket Y \rrbracket$ is better than $\llbracket X \rrbracket \cup \llbracket Y \rrbracket$. Applying Separability again implies that this is true iff $w(x) \cup \llbracket Y \rrbracket$ is better than $\llbracket Y \rrbracket$. So, $w(x)$ is above the critical set in Y . Since Y is an arbitrary population, $w(x)$ is above the critical set in all populations. The same applies *mutatis mutandis* for welfare levels below the critical set. Any welfare level neither above nor below the critical set must be in it.

population X , it is contributively good (bad/strictly neutral/weakly neutral) relative to all populations. Separability also implies that each life's comparative value is constant across populations. If x is contributively better than (worse than/equally good as/incommensurable with) y relative to some population X , it is contributively better than (worse than/equally good as/incommensurable with) y relative to all populations. Therefore, I drop the relativisation to particular populations in what follows.

A sixth assumption is *Archimedeanism about Populations*:

Archimedeanism about Populations

For any life x , any contributively good (bad) life y , and any number m , there is some number n such that a population of n lives at $w(y)$ is better (worse) than a population of m lives at $w(x)$.⁶

As extant theorems demonstrate, any population axiology satisfying the above principles is a Critical-Set View.⁷ On these views, one or more (possibly transformed) welfare levels are designated as critical. The contributive value of a life x relative to a critical level q is then determined by applying some function f to x 's welfare level and then subtracting q :

$$c(x)_q = f(w(x)) - q$$

In what follows, I take f to be the identity function, but my discussion applies equally to any Critical-Set View on which f is strictly increasing. That includes Critical-Set Views with a prioritarian (strictly increasing and strictly concave) weighing function. Any view in which f is *not* strictly increasing will violate the Principle of Personal Good. Restricting f to strictly increasing functions implies that a life x is personally better than (worse than/equally good as) a life y iff x is contributively better than (worse than/equally good as) y , so I drop the 'personally' and 'contributively' from comparisons between lives in what follows.

The value of a population X relative to a critical level q is the sum of the contributive values of each life x_i in X relative to q :

$$v(X)_q = \sum_i c(x_i)_q$$

And a population X is at least as good as a population Y iff $\sum_i c_i(x)_q \geq \sum_i c_i(y)_q$ relative to each q in the critical set Q . If neither X nor Y is at least as good as the other, they are incommensurable.

⁶ Denying this last assumption yields a Lexical View. For discussion of Lexical Views, see Kitcher (2000), Thomas (2018), Nebel (2021), Carlson (forthcoming), and Thornley (2021; forthcoming).

⁷ See, for example, Blackorby, Bossert, and Donaldson's Theorem 6.10 (2005, chap. 6), Broome (1991, 70; 2004), Thomas (forthcoming), and Bossert's Theorem 3 (forthcoming).

Here is an example to illustrate. Suppose we have two populations, X and Y , represented by the following distributions:

$$\begin{aligned} \llbracket X \rrbracket &= \{5\} \\ \llbracket Y \rrbracket &= \{2, 2, 2\} \end{aligned}$$

On a Critical-Set View with a single critical level at 0, X is worse than Y .⁸ On a view with a single critical level at 4, X is better than Y .⁹ On a view with multiple critical levels including 0 and 4, X is incommensurable with Y , because the value of X is not at least as great as the value of Y relative to $q = 0$ and the value of Y is not at least as great as the value of X relative to $q = 4$.

The six principles prior to this example constitute the common core of Critical-Set Views. The following five choice-points divide the class. First, a Critical-Set View's critical set can comprise either a single critical level or multiple critical levels, forming a critical range.¹⁰ The former are Critical-Level Views and the latter are Critical-Range Views. Critical-Level Views claim that lives at the critical level are contributively *strictly* neutral, so that any population X and an otherwise identical population with an extra life at the critical level are equally good.¹¹ That implies that any two populations are commensurable. Critical-Range Views claim that lives within the critical range are contributively *weakly* neutral, so that any two populations X and X -plus-an-extra-life-within-the-critical-range are incommensurable.

The second choice-point concerns the personally neutral set. This too can comprise either a single personally neutral level or a personally neutral range. Neutral-Level Views claim that lives at the personal neutral level are personally *strictly* neutral, so that they are personally equally good as the standard. *Neutral-Range Views* claim that lives within the personal neutral range are personally *weakly* neutral, so that they are personally incommensurable with the standard. Henceforth, I drop the 'personally' from such expressions. *Neutral* levels and ranges are personally neutral. *Critical* levels and ranges are contributively neutral.

The third choice-point is one on which I have already taken a stand. Critical-Range and Neutral-Range Views can interpret their critical and neutral ranges as ranges of incommensurability (Blackorby, Bossert, and Donaldson 1996), parity (Qizilbash 2007; 2018; Rabinowicz 2009), indeterminacy (Broome 2004), some other value-relation, or any combination of the aforementioned

⁸ $v(X)_0 = (5 - 0) = 5$ and $v(Y)_0 = (2 - 0) + (2 - 0) + (2 - 0) = 6$

⁹ $v(X)_4 = (5 - 4) = 1$ and $v(Y)_4 = (2 - 4) + (2 - 4) + (2 - 4) = -6$

¹⁰ Given the Principle of Personal Good, the critical set cannot feature any gaps. In other words, there cannot be a non-critical level between two critical levels.

¹¹ Critical-Level Views could instead claim that lives at the critical level are contributively *weakly* neutral, so that X and X -plus-an-extra-life-at-the-critical-level are *incommensurable*, but this move has few advantages and I have not seen it made.

phenomena. I adopt the language of incommensurability in this paper, but my discussion can be translated into other terms without significant change to its import.

The fourth choice-point concerns the form of the function f . As mentioned above, I take f to be the identity function, but my discussion applies equally to any view on which f is strictly increasing.

The fifth choice-point concerns the relative positions of the critical and neutral sets. The options available at this stage depend on the directions taken at the first and second choice-points, so I outline them in Figure 1. The numbers at each terminus indicate which of the objections listed below apply to that view.

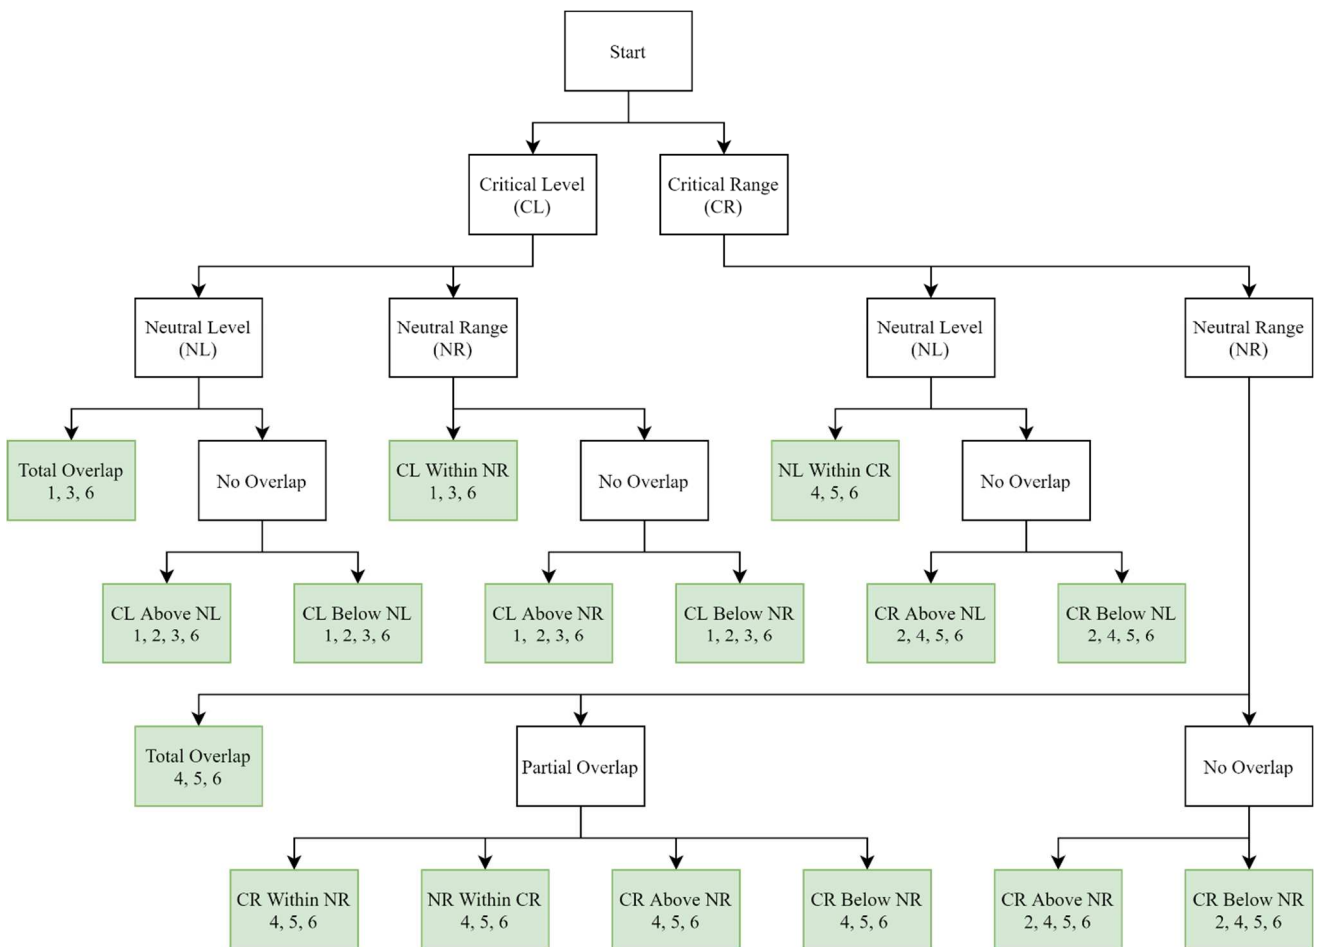


Figure 1

Many of the views in this taxonomy have never been advocated in print, but I lay them all out here for the sake of completeness. Four views that have been defended in print are the *Total View*, a *Positive Critical-Level View*, a *Critical-Range View*, and a *Neutral-Range View*. I diagram them below. Horizontal lines denote that the corresponding welfare level is personally/contributively strictly neutral. Boxes denote that the corresponding welfare levels are personally/contributively weakly neutral. Welfare levels above (below) the

horizontal line or shaded box are personally/contributively good (bad). The numbers are purely illustrative.

First, the Total View:




Figure 2

This view is defended by Hudson (1987), Tännsjö (2002), and Huemer (2008), amongst others. There is a single overlapping neutral level and critical level, so that a life is personally good (bad/strictly neutral) iff it is contributively good (bad/strictly neutral). Any two populations are commensurable.

Second, a Positive Critical-Level View:

A Positive Critical-Level View




Figure 3

Blackorby, Bossert, and Donaldson (2005) and Bossert (forthcoming) defend a Positive Critical-Level View. There is a single critical level above a single neutral level, so some personally good lives are contributively bad. Any two populations are commensurable.

Third, a Critical-Range View:




Figure 4

A view of this kind is defended by Broome (2004), who interprets the critical range as a range of indeterminacy, along with Qizilbash (2007; 2018) and Rabinowicz (2009), who each interpret the critical range as a range of parity. There is a single neutral level but a critical range, so any overlap can be partial at most. Here I present a version of the view in which the neutral level corresponds to the lowest welfare level in the critical range. On Critical-Range Views, some pairs of populations are incommensurable.

Finally, a Neutral-Range View:




Figure 5

Rabinowicz discusses a view of this kind in more recent work (forthcoming) and Gustafsson (2020) defends a view of this form in which there is a neutral and critical range for temporal welfare levels as well as lifetime welfare levels. On Neutral-Range Views, there is a neutral range and critical range that totally overlap, so a life is personally good (bad/weakly neutral) iff it is contributively good (bad/weakly neutral). Some pairs of populations are incommensurable.

3. Objections to Critical-Set Views

Many varieties of Critical-Set View are subject to the same objections. At least three of the following six objections apply to each view in the taxonomy.

1. Maximal Repugnance

Any Critical-Set View on which lives barely worth living are contributively good will imply the Repugnant Conclusion:

Each population of wonderful lives is worse than some population of lives barely worth living. (Parfit 1984, 388)

And any Critical-Set View on which lives barely worth *not* living are contributively bad will imply the Mirrored Repugnant Conclusion:

Each population of awful lives is better than some population of lives barely worth not living. (Carlson 1998, 297; Broome 2004, 213; Gustafsson 2020, 85)

Both of these consequences arise because Critical-Set Views accept Archimedeanism about Populations: a population of enough contributively good (bad) lives can be better (worse) than any other population.

However, as Broome (2004, 213) and Rabinowicz (2009, 406; forthcoming, 79) note, the repugnance of these conclusions is attenuated if a wide range of lives are personally neutral. In that case, lives barely worth living are much better than lives barely worth not living. What makes the Repugnant Conclusion and its mirror troubling is the presumed similarity of lives barely worth living and lives barely worth not living. With that in mind, I define *Maximal Repugnance* as follows:

There is a life x and a life y that is identical but for an extra two hangnail's worth of pain such that (1) each population of wonderful lives is worse than some population of x lives and (2) each population of awful lives is better than some population of y lives.

Note that I drop the specification that x is barely worth living and y is barely worth not living. This feature is not necessary for repugnance. Suppose, for example, that we accept a view which implies Maximal Repugnance for a life x that is significantly personally good. This move mitigates the force of implication (1): we might be quite happy to accept that each population of wonderful lives is worse than some population of significantly personally good lives. But it exacerbates the implausibility of implication (2): if x is significantly personally good, then y is personally good, and it is hard to believe that each population of awful lives is better than some population of personally good lives. More generally, at least one of implications (1) and (2) will be implausible no matter how good x and y are.

Given that an extra two hangnail's worth of pain suffice to make a life two welfare levels worse, any view with just one critical level will have this consequence. In short, all Critical-Level Views imply Maximal Repugnance.

2. Sadism

This objection requires me to introduce a new distinction. If welfare levels are *dense*, then for any two distinct welfare levels, there is a welfare level between them. If welfare levels are *not dense*, then there is at least one pair of distinct welfare levels with no welfare level in between (Arrhenius 2011, 6).

Given denseness, any view on which there is no overlap between the critical set and the neutral set implies a Sadistic Conclusion. If the critical set is above the neutral set, the view implies the original Sadistic Conclusion:

Each population of awful lives is better than some population of personally good lives. (Arrhenius 2000a, 256)

That is because positioning the critical set above the neutral set entails that some personally good lives are contributively bad. Given Archimedeanism, each population is worse than a population of enough contributively bad lives.

If the critical is below the neutral set, the view implies the Mirrored Sadistic Conclusion:

Each population of wonderful lives is worse than some population of personally bad lives. (Gustafsson 2020, 85)

That is because positioning the critical set below the neutral set entails that some personally bad lives are contributively good. The Archimedean principle does the rest.

If denseness is false, we could endorse a Critical-Set View on which there is no overlap between the critical set and the neutral set and yet no welfare level between the two sets. These kinds of views imply only weaker forms of sadism. If the bottom of the critical set is one welfare level above the top of the neutral set, the view implies a Weaker Sadistic Conclusion:

Each population of awful lives is better than some population of personally neutral lives.

If the top of the critical set is one welfare level below the bottom of the neutral set, the view implies a Weaker Mirrored Sadistic Conclusion:

Each population of wonderful lives is worse than some population of personally neutral lives.

These conclusions are more plausible than the pair above, but that is faint praise. In fact, comparison with the previous subsection will show that they could equally be called Stronger Mirrored and Stronger Repugnant Conclusions, respectively.¹²

All views with no overlap between the critical set and the neutral set imply some form of sadism.

3. Strong Superiority Across Slight Differences

Consider a sequence of lives beginning with a contributively good life x_1 . We reach x_2 by making x_1 slightly worse. Perhaps x_2 is identical to x_1 but for one

¹² I use the words ‘weaker’ and ‘stronger’ rather than ‘weak’ and ‘strong’ to distinguish these conclusions from the Weak Sadistic Conclusion and Strong Repugnant Conclusion that appear in Gustafsson (2020, 86) and Meacham (2012, 270) respectively.

extra hangnail’s worth of pain. We reach x_3 by making x_2 slightly worse, and so on, until after a finite number of slight detriments we reach x_n , a contributively bad life.

On Critical-Level Views, each life is either contributively good, contributively strictly neutral, or contributively bad. That means that, in our sequence, there is some contributively good life x_k such that x_{k+1} is either contributively strictly neutral or contributively bad. That in turn implies that x_k has positive contributive value, while x_{k+1} ’s contributive value is non-positive. Adding positive numbers can never yield a non-positive number, and vice versa, so Critical-Level Views imply that any number of lives x_k is better than any number of lives x_{k+1} . Call this implication *Strong Superiority Across Slight Differences* (SSASD).¹³ Two corollaries of SSASD bring out its implausibility: a single life x_k is better than any number of lives x_{k+1} , and a single life x_{k+1} is worse than any number of lives x_k .

We might claim that this implication is of little concern: x_k is contributively good and x_{k+1} is not, so the strong superiority of x_k over x_{k+1} should come as no surprise. But this level of description masks the difficulty. The life x_k might be long and turbulent, featuring soaring highs and crushing lows. Amidst these peaks and troughs, we might expect a hangnail to pale almost into axiological insignificance. But Critical-Level Views imply that this drop in the ocean makes all the difference, so that any number of lives without the hangnail is better than any number of lives with it.

4. Strong Noninferiority Across Slight Differences

SSASD might spur us to adopt a Critical-Range View. On Critical-Range Views, a range of lives are contributively weakly neutral. If this range is wide enough, our x -sequence will contain no lives x_k and x_{k+1} such that x_k is contributively good and x_{k+1} is contributively strictly neutral or bad. If x_k is the last contributively good life in the sequence, then x_{k+1} will be contributively *weakly* neutral. That means that Critical-Range Views can avoid SSASD, because it is not the case that any number of contributively good lives is better than any number of contributively weakly neutral lives. Instead, a population of enough contributively weakly neutral lives is incommensurable with each population of contributively good lives, as the following example shows.

Suppose that all welfare levels q between 0 and 4 inclusive are in the critical set Q . And suppose that $w(x_k) = 4.01$ and $w(x_{k+1}) = 3.99$. Population X consisting of a single life x_k is better than population Y consisting of a single life x_{k+1} , because $v(X) > v(Y)$ for each q in Q . But X is incommensurable with

¹³ For discussions of superiority and noninferiority in axiology, see Arrhenius and Rabinowicz (2015b), Nebel (2021), and Thornley (forthcoming).

population Z consisting of two lives x_{k+1} . X has greater value than Z relative to $q = 4$,¹⁴ but Z has greater value than X relative to $q = 0$.¹⁵

More generally, each contributively weakly neutral life has positive value relative to some q .¹⁶ That implies that each population has less value than a population of enough contributively weakly neutral lives relative to that q . Therefore, each population is not better than a population of enough contributively weakly neutral lives.

However, Critical-Range Views still imply *Strong Noninferiority Across Slight Differences*: for some x_k and x_{k+1} in our x -sequence, any number of lives x_k is *not worse than* any number of lives x_{k+1} . We can see this with the above example. No matter how many lives x_k are contained in X , and no matter how many lives x_{k+1} are contained in Z , X will have greater value than Z relative to $q = 4$. So, X is not worse than Z , no matter what their respective sizes. Of course, this implication is less troubling than Critical-Level Views' *Strong Superiority Across Slight Differences*. But in some sequences, at least one of the discontinuities implied by Critical-Range Views must occur in a more counterintuitive place, as I demonstrate below.

Consider a new sequence. Each life in this sequence features a blank period, free of any good or bad components. We can imagine it as a minute of dreamless sleep. The first life in the sequence y_0 also features a period of constant happiness of length n hours, and nothing else. The second life y_1 is identical, except that the happiness lasts $n - 1$ hours. y_2 's happiness lasts $n - 2$ hours, and so on. Call all such lives featuring only good and neutral components 'straightforwardly-better-than-blank.' y_n features only the blank period, and so qualifies as a *blank life*, featuring no good or bad components whatsoever (Broome 2004, 208). y_{n+1} features the blank period plus one hour of suffering, y_{n+2} features the blank period plus two hours of suffering, and so on. The last life in the sequence is y_{2n} , featuring the blank period plus n hours of suffering. Call all such lives featuring only bad and neutral components 'straightforwardly-worse-than-blank.'

Intuitively, the first discontinuity in this sequence occurs between y_{n-1} and y_n . That is, y_{n-1} is strongly noninferior to y_n : any number of lives y_{n-1} featuring

¹⁴ $v(X)_4 = (4.01 - 4) = 0.01$ and $v(Z)_4 = (3.99 - 4) + (3.99 - 4) = -0.02$

¹⁵ $v(X)_0 = (4.01 - 0) = 4.01$ and $v(Z)_0 = (3.99 - 0) + (3.99 - 0) = 7.98$

¹⁶ We might think that lives with 0 welfare are a counterexample to this claim. They do not have positive value relative to any q between 0 and 4 inclusive. But these lives are not contributively weakly neutral. On our definitions, they are contributively bad. Here is why. Suppose $w(x) = 0$. Then, for any X , the value of $\llbracket X \rrbracket$ is at least as great as the value of $\llbracket X \rrbracket \cup w(x)$ relative to each q , so $\llbracket X \rrbracket$ is at least as good as $\llbracket X \rrbracket \cup w(x)$. But the value of $\llbracket X \rrbracket \cup w(x)$ is not at least as great as the value of $\llbracket X \rrbracket$ relative to each q , so $\llbracket X \rrbracket \cup w(x)$ is not at least as good as $\llbracket X \rrbracket$. Therefore, $\llbracket X \rrbracket \cup w(x)$ is worse than $\llbracket X \rrbracket$, and x is contributively bad. This is strange because $w(x)$ is in the critical range, but this strangeness turns out to be of little consequence. We just need to bear in mind that only lives within the boundaries of the critical range are contributively weakly neutral.

1 hour of happiness is not worse than any number of blank lives y_n . And, again intuitively, the second discontinuity in this sequence occurs between y_n and y_{n+1} . That is, y_{n+1} is strongly *nonsuperior* to y_n : any number of lives y_{n+1} featuring 1 hour of suffering is *not better* than any number of blank lives y_n . These two claims remain intuitive when we replace ‘hours’ with ‘minutes,’ ‘seconds,’ ‘milliseconds,’ and so on.

But Critical-Range Views must deny at least one of these claims. Recall that, on Critical-Range Views, more than one welfare level is contributively neutral. Therefore, in any sequence with sufficiently small welfare differences between adjacent lives, more than one life is contributively weakly neutral. We can make the welfare differences between adjacent lives in our y -sequence arbitrarily small by replacing ‘hours’ with smaller units of time, so for some such unit, more than one life in our y -sequence is contributively weakly neutral.

Suppose for illustration that, when the unit of time is seconds, y_{n-1} and y_n are the contributively weakly neutral lives. In that case, y_{n-2} , the last contributively good life, is strongly noninferior to y_{n-1} , the first contributively weakly neutral life. In other words, a population of any number of lives featuring two seconds of happiness is not worse than a population of any number of lives featuring one second of happiness. That implies that a population of a *single* life featuring two seconds of happiness is not worse than a population of any number of lives featuring one second of happiness. But this implication seems absurd. The only difference between these lives is the duration of happiness, the latter population can feature an arbitrarily longer duration of happiness than the former, and yet the latter population can never be better than the former.

We get a mirror of this implication if we instead suppose that y_n and y_{n+1} are the contributively weakly neutral lives. In that case, a population of any number of lives featuring two seconds of suffering is not better than a population of any number of lives featuring one second of suffering. Though this latter population can feature an arbitrarily long duration of suffering, it can never be worse than a population of a single life featuring just two seconds of suffering. This too seems absurd.

Nothing hinges on the particular lives chosen to illustrate this dynamic. Any Critical-Range View will imply that (1) a population consisting of a single straightforwardly-better-than-blank life is not worse than a population consisting of any number of straightforwardly-better-than-blank lives identical but for a slightly smaller quantity of good, or (2) a population consisting of a single straightforwardly-worse-than-blank life is not better than a population consisting of any number of straightforwardly-worse-than-blank lives identical but for a slightly smaller quantity of bad.

5. Maximal Greediness

Critical-Range Views face another difficulty. As Broome points out (2004, 169–70, 202–5), they imply that contributively weakly neutral lives can ‘swallow up’ and neutralise goodness and badness. Here is an illustration of what that means. Suppose again that all welfare levels between 0 and 4 inclusive are contributively neutral. And suppose that population A consists of a single life x at welfare level 20. We reach population B by making two changes. We reduce x ’s welfare by 1 and add a life y at welfare level 2. The combined effect of these changes might seem bad. We made one person worse-off and added a life that is contributively neutral. But our Critical-Range View implies that these changes are not bad. Neither A nor B ’s value is at least as great as the other relative to each q in Q , so the two populations are incommensurable.¹⁷ Our Critical-Range View also implies that A is incommensurable with C – in which x ’s welfare is 18 and there are two lives at welfare level 2 – and D – in which x ’s welfare is 17 and there are three lives at welfare level 2 – and so on. This process can continue indefinitely. A will also be incommensurable with a population Z , in which x ’s welfare is extremely negative and there is some large number of contributively neutral lives.

Broome and I find this ‘greedy neutrality’ concerning, but others are happy to bite the bullet (Rabinowicz 2009; Frick 2017; Gustafsson 2020). In any case, the worry can be sharpened.

Note first that the size of population A need not be restricted to a single life: adding enough contributively weakly neutral lives can neutralise any finite loss of welfare for existing people. And suppose that blank lives are contributively weakly neutral. In that case, for any arbitrarily good population and any arbitrarily bad population, there is some population of blank lives – featuring no good or bad components whatsoever – such that the good population plus the blank lives is not better than the bad population. This implication seems difficult to accept.

It gets worse. Consider again our y -sequence above. Given that the unit of time is sufficiently small, Critical-Range Views imply that more than one life in this sequence is contributively weakly neutral. For illustration, suppose that the blank life y_n and the straightforwardly-better-than-blank life y_{n-1} are contributively weakly neutral. In that case, we can replace ‘blank lives’ with ‘straightforwardly-better-than-blank lives’ in the above paragraph. For any arbitrarily good population and any arbitrarily bad population, there is some population of straightforwardly-better-than-blank lives – featuring no bad components whatsoever and some happiness – such that the good population plus the straightforwardly-better-than-blank lives is not better than the bad

¹⁷ On $q = 4$, $v(A)_4 = (20 - 4) = 16$ and $v(B)_4 = (19 - 4) + (2 - 4) = 13$. On $q = 0$, $v(A)_0 = (20 - 0) = 20$ and $v(B)_0 = (19 - 0) + (2 - 0) = 21$.

population. The former population features only neutral and good components, the latter population might feature only bad components, and yet this Critical-Range View implies that the former is not better than the latter.

If the straightforwardly-worse-than-blank life y_{n+1} is contributively weakly neutral, we get a mirror of this implication. For any arbitrarily good population and any arbitrarily bad population, there is some population of straightforwardly-worse-than-blank lives – featuring no good components whatsoever and some suffering – such that the bad population plus the straightforwardly-worse-than-blank lives is not worse than the good population. Call implications of this kind *Maximal Greediness*.

Shifting the critical range away from blank lives fails to mitigate the difficulty. If the critical range is above or below the welfare level of a blank life, then some other life in our y -sequence will be contributively weakly neutral. No matter where the critical range is placed, we get Maximal Greediness.

6. No Incommensurability Between Lives or Same-Number Populations

On Critical-Level Views, a population's value can be represented by a real number. Since any two real numbers are commensurable (a is at least as great as b or b is at least as great as a), Critical-Level Views imply that any two populations are commensurable: X is at least as good as Y or Y is at least as good as X .

However, universal commensurability seems implausible. Consider the following Small Improvement Argument (De Sousa 1974; Chang 2002). Suppose that X consists of 10 wonderful lives and Y consists of 100 very good lives. Neither X nor Y is better than the other.¹⁸ If any two populations are commensurable, X and Y are equally good. But if X and Y are equally good, then any population better than Y is better than X . Y^+ , consisting of 101 very good lives, is better than Y but not better than X . Therefore, X and Y are not equally good. They are incommensurable.

Critical-Range Views can account for this incommensurability. They can claim that X has greater value than Y relative to one level in the critical range and that Y has greater value than X relative to another level. But this explanation cannot account for all plausible instances of incommensurability. In particular, it cannot account for the incommensurability of same-number populations.

This is easiest to see in the single life case. Critical-Set Views assume that a life's welfare can be represented by a real number. Since any two real numbers are commensurable, this assumption implies that any two lives are commensurable: x is at least as good as y or y is at least as good as x .

¹⁸ Those who disagree should tweak the numbers and/or adjectives.

Now recall Critical-Set Views' equation for the value of a population X relative to a neutral level q :

$$v(X)_q = \sum_i (f(w(x_i)) - q)$$

Since this equation is a sum of transformed welfare levels, assuming that a life's welfare can be represented by a real number implies that a population's value relative to a neutral level can be represented by a real number. That in turn implies that the value of any two populations relative to a neutral level is commensurable. Formally,

- (1) For any populations X and Y and any contributively neutral level q , $v(X)_q \geq v(Y)_q$ or $v(Y)_q \geq v(X)_q$.

Now let X and Y stand for arbitrary same-number populations and q stand for an arbitrary neutral level such that $v(X)_q \geq v(Y)_q$. Substituting in the equations for $v(X)_q$ and $v(Y)_q$ gives us the following inequality:

$$\sum_i (f(w(x_i)) - q) \geq \sum_i (f(w(y_i)) - q)$$

This inequality can also be expressed as follows, with n representing the size of populations X and Y :

$$\left(\sum_i f(w(x_i)) \right) - nq \geq \left(\sum_i f(w(y_i)) \right) - nq$$

The terms involving q can then be cancelled from each side:

$$\sum_i f(w(x_i)) \geq \sum_i f(w(y_i))$$

Therefore, the inequality is true for all values of q , and X is at least as good as Y . Since X , Y , and q were arbitrary, we can conclude:

- (2) For any same-number populations X and Y and any contributively neutral level q , if $v(X)_q \geq v(Y)_q$, then X is at least as good as Y .

Together, (1) and (2) imply:

- (3) For any same-number populations X and Y , X is at least as good as Y or Y is at least as good as X .

In sum, Critical-Set Views imply that any two same-number populations are commensurable.

However, universal commensurability of same-number populations seems implausible. Consider another Small Improvement Argument. Suppose that x is a turbulent life, featuring soaring highs and crushing lows, and that y is a drab life, featuring only Muzak and potatoes. If we fix the relative quantities of x 's

highs and lows in the right way, neither x nor y is better than the other. Yet x and y cannot be equally good, because a slightly less drab life y^+ – featuring Muzak, potatoes, and ketchup – is better than y but not better than x . Therefore, x and y are incommensurable. Similar arguments establish the incommensurability of same-number populations of other sizes.

Partly on the basis of such arguments, advocates of Critical-Set Views have started to incorporate incommensurability and indeterminacy into their orderings of lives. Broome (forthcoming), for example, states that some pairs of lives are obviously indeterminately related, but offers no explanation for why this is so. Rabinowicz (forthcoming), meanwhile, offers a fitting-attitudes analysis of parity – one species of incommensurability – according to which two lives are on a par iff it is permissible to prefer either life to the other. And Gustafsson (2020) accounts for incommensurability between lives by claiming that there is a neutral range of temporal welfare levels. Adding a moment within this range to a life renders the new life incommensurable with the original.

Gustafsson’s move strikes me as a step in the right direction. However, his view cannot account for the incommensurability between same-length lives, for the same reason that Critical-Range Views cannot account for the incommensurability between same-number populations. Gustafsson might claim that any two lives of the same length are commensurable, but this claim seems implausible. The Small Improvement Argument involving drab and turbulent lives remains convincing if we specify that the lives are the same length.

Rabinowicz’s (forthcoming) account is incomplete but, I believe, more promising. He claims that ‘life wellbeing is a many-dimensional concept,’ that ‘specifying its level requires characterizing a life with respect to several relevant dimensions,’ and that ‘different weight assignments’ to these relevant dimensions give rise to incommensurability between lives (forthcoming, 81). This notion of ‘different weight assignments’ forms the core of the Imprecise Exchange Rates View.

4. Imprecise Exchange Rates

In some cases, undergoing a bad for the sake of some good is *worth it*. For example, it is worth going to the dentist to prevent tooth decay. The good of having healthy teeth outweighs the bad of the trip. Other trade-offs are *not worth it*. Getting up at 4am and walking to work to save the £2 bus fare is not worth it. The bad outweighs the good. In still other cases, undergoing a bad for the sake of some good is *neither worth it nor not worth it*, as the following Small Improvement Argument shows.

A parent says to their child, ‘No dessert unless you finish your dinner.’ The child knows exactly what finishing dinner involves. They are all-too-familiar with

the taste of Brussels sprouts and can see five of them left on the plate. They also know what dessert will be like. The jelly is sitting on the counter and promises to taste as good as it always has. In this case, eating the sprouts may be neither worth it nor not worth it. And a small improvement to the child’s predicament need not resolve the issue. Suppose that the parent takes pity on the child and removes one sprout from the plate. That need not ensure that finishing dinner is now worth it.

I claim that cases of this kind are evidence that various *exchange rates* – between pairs of goods, between pairs of bads, and between goods and bads – are imprecise. This imprecision renders certain goods incommensurable with other goods, certain bads incommensurable with other bads, and certain combinations of goods and bads incommensurable with other combinations. In the child’s case, eating both the sprouts and the jelly is incommensurable with eating neither. This incommensurability between goods, bads, and their combinations is the source of incommensurability between lives. The child’s life in which they eat the sprouts and jelly is incommensurable with the otherwise identical life in which they eat neither.

That constitutes the motivation for the Imprecise Exchange Rates (IER) View. Now for the formalisation. Recall that Critical-Set Views begin with an ordering of lifetime welfare levels. The IER-View begins instead with a set of orderings: one for each dimension of good and bad within a life. The exact form of the view thus depends on our theory of welfare. If we accept the simplest hedonist theory, there are just two orderings: one of happiness and one of suffering. If we accept an objective list theory, there are more orderings: perhaps one of love, one of virtue, one of false belief, etc. Welfare levels are thus given by vectors. Suppose, for example, that we accept an objective list theory on which happiness, love, suffering, and false belief comprise the dimensions of good and bad. Then the welfare of a life x can be expressed as follows:

$$w(x) = \langle h(x), l(x), s(x), f(x) \rangle$$

I assume that each dimension is representable by a real-valued function. I also assume that the values of each function are interpersonally comparable (so that we can make claims like ‘Ada’s life as an artist features more happiness than Bob’s life as a baker.’) and measurable on a ratio-scale (so that we can make claims like ‘Ada’s life as an artist features twice the suffering of Ada’s life as a baker.’). Blank lives – featuring no good or bad components whatsoever – score 0 on each dimension.

Each ratio-scale is independent, so we cannot yet compare values across dimensions. We cannot make claims like ‘In Ada’s life as an artist, her happiness outweighs her suffering.’ Comparisons of this kind are only possible given a specified *proto-exchange-rate* r : a vector of two or more real numbers strictly

greater than 0 and summing to 1, denoting the relative weight granted to each dimension of good and bad. On the objective list theory above, for example, each proto-exchange-rate r will take the form $\langle r_h, r_l, r_s, r_f \rangle$, where r_h denotes the weight granted to happiness, r_l denotes the weight granted to love, and so on. Letting x represent Ada's life as an artist, the claim that her happiness outweighs her suffering on a given r will be true iff $r_h h(x) > r_s s(x)$.

On the IER-View, only welfare levels *on a given r* are expressible as a real number. Continuing with our example objective list theory, the equation is as follows:

$$w(x)_r = r_h h(x) + r_l l(x) - r_s s(x) - r_f f(x)$$

The value of a population on r is the sum of the welfare levels of each of its lives on r :

$$v(X)_r = \sum_i w(x_i)_r$$

We then account for incommensurability by claiming that there are multiple proto-exchange-rates r in the set of all proto-exchange-rates R . A life x is at least as good as a life y iff $w(x)_r \geq w(y)_r$ on each r in R . And a population X is at least as good as a population Y iff $\sum_i w(x_i)_r \geq \sum_i w(y_i)_r$ on each r in R .¹⁹

In what follows, I mostly discuss a simple hedonist version of the IER-View, in which the welfare level of a life x is represented by a vector of happiness and suffering, $\langle h(x), s(x) \rangle$, with the functions h and s normalised so that the proto-exchange-rate r composed of $r_h = 0.5$ and $r_s = 0.5$ falls within the set R . In other words, undergoing an episode of suffering represented by a number k for the sake of an episode of happiness represented by that same number k is neither worth it nor not worth it. I adopt hedonism purely for the sake of simplicity. Its two dimensions are sufficient to illustrate the most important advantages and drawbacks of the IER-View. My discussion below applies equally to variants of the view with more dimensions.

5. Advantages of the Imprecise Exchange Rates View

The IER-View has several advantages over the Critical-Set Views discussed above. Here are four.

¹⁹ Rabinowicz (forthcoming, 83–84) offers a similar formalisation. His formalisation, however, takes a set of permissible preferential ratio-scales over the set of lives as primitive. It does not specify how the dimensions of welfare weigh against each other.

1. Some Incommensurability Between Lives and Same-Number Populations

The first advantage is that the IER-View offers a simple and plausible account of incommensurability between lives and same-number populations. Recall that a life is at least as good as another iff its welfare is at least as great on each r in R . If R contains more than one r , then some pairs of lives are incommensurable: neither is at least as good as the other.

Consider an example. Suppose that R contains each r in which $0.4 \leq r_h \leq 0.6$. Since $r_h + r_s = 1$, $r_s = 1 - r_h$. In that case, life x – at welfare level $\langle 4, 1 \rangle$ – is incommensurable with life y – at welfare level $\langle 10, 6 \rangle$. The value of x is greater on $r_h = 0.4$,²⁰ but the value of y is greater on $r_h = 0.6$.²¹ This is as it should be. Taking on the extra suffering in y for the sake of the extra happiness is neither worth it nor not worth it.

The IER-View also gives us the right result in Small Improvement Cases. A slightly improved life $y^+ = \langle 10 + e, 6 \rangle$ comes out better than y and incommensurable with x . That is because the IER-View accounts for incommensurability between lives while respecting a certain kind of dominance:

Dominance over Dimensions

For any lives x and y and any set of proto-exchange-rates R , if for each good dimension g , x features at least as much g as y , and for each bad dimension b , x features at most as much b as y , x is at least as good as y . If, in addition, x features more g than y for some g or less b than y for some b , x is better than y .²²

Another implication is related. Let us say that a pair of proto-exchange-rates *differ in optimism* iff they differ in the total weight granted to all dimensions of good taken together.²³ The implication is that if R contains proto-exchange-rates

²⁰ $w(x)_{r_h=0.4} = 0.4 \times 4 - 0.6 \times 1 = 1$ and $w(y)_{r_h=0.4} = 0.4 \times 10 - 0.6 \times 6 = 0.4$

²¹ $w(x)_{r_h=0.6} = 0.6 \times 4 - 0.4 \times 1 = 2$ and $w(y)_{r_h=0.6} = 0.6 \times 10 - 0.4 \times 6 = 3.6$

²² Here is a sketch of the proof. x is at least as good as y on any R iff $r_h h(x) - r_s s(x) \geq r_h h(y) - r_s s(y)$ for any $0 < r_h < 1$ and $r_s = 1 - r_h$. Rearranging this equation gives $r_h(h(x) - h(y)) + r_s(s(y) - s(x)) \geq 0$. If x dominates y , then $h(x) \geq h(y)$ and $s(y) \geq s(x)$, so each term on the left-hand-side of the inequality in the previous sentence is non-negative. Therefore, the weak inequality holds. If, in addition, x features more happiness or less suffering than y , then at least one term on the left-hand-side of the inequality is positive, so the strict inequality holds. This proof can be extended to any number of dimensions of good and bad.

²³ Here is an example. Return briefly to our objective list theory in which happiness, love, suffering, and false belief are the dimensions of good and bad, and consider the following three proto-exchange-rates: $r_1 = \langle 0.3, 0.2, 0.1, 0.4 \rangle$, $r_2 = \langle 0.2, 0.3, 0.1, 0.4 \rangle$, and $r_3 = \langle 0.3, 0.3, 0.1, 0.3 \rangle$. r_1 and r_2 are distinct, because r_1 assigns more weight to happiness while r_2 assigns more weight to love. But they are equally optimistic, because they both assign a weight of 0.5 to both dimensions

that differ in optimism, then only lives featuring identical total quantities of good and bad can be equally good.²⁴ That means that lives such as $\langle 4, 4 \rangle$ and $\langle 5, 5 \rangle$ come out incommensurable on the IER-View. This result is exactly what we want. Undergoing the extra suffering for the sake of the extra happiness is neither worth it nor not worth it. If $\langle 4, 4 \rangle$ and $\langle 5, 5 \rangle$ were judged equally good, the view would generate counterintuitive verdicts in small improvement cases. For example, $\langle 4, 4 \rangle$ would be worse than $\langle 5, 5 - e \rangle$, for any $e > 0$. Henceforth, I assume that R contains proto-exchange-rates that differ in optimism.

The above three points are true of populations as well as lives. If R contains more than one r , then some pairs of populations (including same-number populations) are incommensurable. If one population weakly (strictly) dominates another over dimensions, then it is at least as good (better). And if R contains proto-exchange-rates that differ in optimism, then only populations featuring identical total quantities of good and bad can be equally good.

2. No Sadism

Recall that Critical-Set Views positing no overlap between the critical set and the neutral set imply some Sadistic Conclusion: either each population of awful lives is better than some population of lives that are not personally bad, or each population of wonderful lives is worse than some population of lives that are not personally good.

The IER-View can avoid this drawback. More precisely, the IER-View avoids sadism if we make the plausible claim that blank lives are personally strictly neutral. This claim implies that *only* blank lives are personally strictly neutral since, as we saw in the last subsection, no lives differing in their quantities of good and bad can be equally good. The extension of personal strict neutrality

of good taken together. r_3 , meanwhile, differs in optimism from both r_1 and r_2 , because r_3 assigns a weight of 0.6 to both dimensions of good taken together.

²⁴ To see this result, note first that equally good lives must have the same value on each proto-exchange-rate. If x has greater value than y on some proto-exchange-rate, y is not at least as good as x , and so the pair cannot be equally good. Now let $g(x)$ denote the total quantity of good in x , $b(x)$ denote the total quantity of bad, and so on, and let r_1 and r_2 denote the total weight assigned to dimensions of good on proto-exchange-rates that differ in optimism. If x and y are equally good, then $r_1 g(x) - (1 - r_1) b(x) = r_1 g(y) - (1 - r_1) b(y)$ and *mutatis mutandis* for r_2 . Rearranging these equations gives $r_1 (g(x) - g(y) + b(x) - b(y)) + b(x) - b(y) = 0$ and *mutatis mutandis* for r_2 . Since both expressions equal 0, they equal each other. Cancelling $b(x) - b(y)$ from each side gives $r_1 (g(x) - g(y) + b(x) - b(y)) = r_2 (g(x) - g(y) + b(x) - b(y))$. Since $r_1 \neq r_2$, the expression $g(x) - g(y) + b(x) - b(y)$ must equal 0. That is true iff $g(x) - g(y) = k$ and $b(x) - b(y) = -k$. If $k > 0$, then $g(x) > g(y)$ and $b(x) < b(y)$. In that case, x is better than y by strict dominance, so they cannot be equally good. If $k < 0$, then y is better than x by strict dominance. The only remaining possibility is that $k = 0$, in which case $g(x) = g(y)$ and $b(x) = b(y)$. Therefore, x and y are equally good only if they feature identical quantities of good and bad.

then matches the extension of contributive strict neutrality since, on the IER-View, only blank lives are contributively strictly neutral. Adding any other kind of life changes the quantity of good and bad in the population, and no populations differing in their quantities of good and bad can be equally good.

This coincidence of personal and contributive strict neutrality is sufficient to establish the coincidence of all personal and contributive value. That is because the IER-View then determines each life's personal and contributive value in the same way: its value is compared to the value of a blank life on each proto-exchange-rate in R . That implies that a life is personally good (bad/strictly neutral/weakly neutral) iff it is contributively good (bad/strictly neutral/weakly neutral). Therefore, the IER-View avoids all instances of sadism.

With the coincidence of personal and contributive value on the IER-View established, I often drop the words 'personal' and 'contributive' in what follows. I graph these coincident values for lives at different welfare levels in Figure 6, on the IER-View with $0.4 \leq r_h \leq 0.6$. A life is good (bad/weakly neutral) iff the point picked out by its quantity of suffering on the horizontal axis and its quantity of happiness on the vertical axis falls within the green (red/white) region. Lives at the origin are blank and hence strictly neutral.

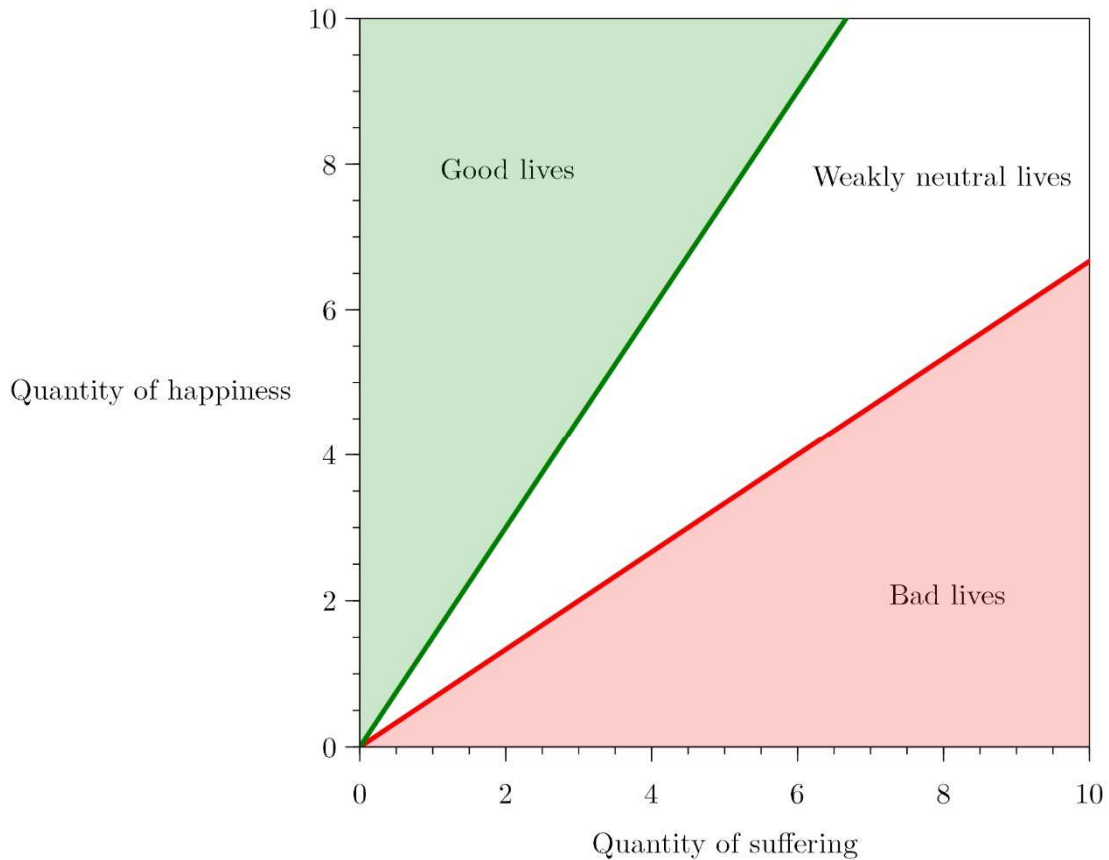


Figure 6

3. Less Concerning Superiority and Noninferiority

As we saw above, Critical-Level Views imply Strong *Superiority* Across Slight Differences (SSASD) in our x -sequence: any number of long, turbulent lives x_k is *better than* any number of lives x_{k+1} identical but for an extra hangnail. Critical-Range Views, meanwhile, imply only Strong *Noninferiority* Across Slight Differences: any number of long, turbulent lives x_k is *not worse than* any number of lives x_{k+1} identical but for an extra hangnail. But on Critical-Range Views, at least one discontinuity of this kind must occur in a counterintuitive place in our y -sequence, so that a single life featuring only neutral components and happiness is not worse than any number of lives each featuring a slightly shorter duration of happiness, or a single life featuring only neutral components and suffering is not better than any number of lives each featuring a slightly shorter duration of suffering.

The IER-View avoids both of these problems. Consider first SSASD. Suppose, for illustration, that an extra hangnail adds 0.02 to a life's quantity of suffering. Suppose also that a turbulent life x_k has welfare $\langle 9, 9 \rangle$. Life x_{k+1} then has welfare $\langle 9, 9.02 \rangle$. Since x_k dominates x_{k+1} , population X consisting of a single life x_k is better than population Y consisting of a single life x_{k+1} . But X is incommensurable with a population Z , consisting of two lives x_{k+1} . X has greater value than Z on $r_h = 0.4$,²⁵ but Z has greater value than X on $r_h = 0.6$.²⁶

We get the same result with lives at many other welfare levels. In fact, the IER-View avoids SSASD in all but a small minority of cases. To see those cases in which SSASD is implied, let $\langle h(x_k), s(x_k) \rangle$ and $\langle h(x_k), s(x_k) + 0.02 \rangle$ represent the welfare levels of x_k and x_{k+1} respectively. x_k is strongly superior to x_{k+1} iff x_k is good and x_{k+1} is strictly neutral or bad, or x_k is strictly neutral and x_{k+1} is bad. Therefore, x_k 's welfare must be non-negative and x_{k+1} 's welfare must be non-positive on each r in R .²⁷ This is so iff x_k 's welfare is non-negative on the most pessimistic proto-exchange-rate $r_h = 0.4$ and x_{k+1} 's welfare is non-positive on the most optimistic proto-exchange-rate $r_h = 0.6$. That yields two inequalities: $0.4h(x_k) - 0.6s(x_k) \geq 0$ and $0.6h(x_k) - 0.4(s(x_k) + 0.02) \leq 0$. Plotting these two inequalities on a graph gives us the following region:

²⁵ $v(X)_{r_h=0.4} = 0.4 \times 9 - 0.6 \times 9 = -1.8$ and $v(Z)_{r_h=0.4} = (0.4 \times 9 - 0.6 \times 9.02) + (0.4 \times 9 - 0.6 \times 9.02) = -3.624$

²⁶ $v(X)_{r_h=0.6} = 0.6 \times 9 - 0.4 \times 9 = 1.8$ and $v(Z)_{r_h=0.6} = (0.6 \times 9 - 0.4 \times 9.02) + (0.6 \times 9 - 0.4 \times 9.02) = 3.584$

²⁷ And at least one life's value must be non-zero on some r in R . In other words, at least one of x_k and x_{k+1} must not be blank. The hangnail's worth of pain ensures that this condition is met.




Figure 7

A life x_k is strongly superior to an otherwise identical life x_{k+1} with an extra hangnail iff the point picked out by $s(x_k)$ on the horizontal axis and $h(x_k)$ on the vertical axis lies within the unshaded region. This is a welcome result. As we can see, an extra hangnail triggers strong superiority only when added to lives featuring very small quantities of happiness and suffering. The IER-View thus gives hangnails their proper axiological due. In blank and nearly-blank lives, they make a difference. In turbulent lives, they pale almost into axiological insignificance.²⁸

I write ‘almost’ because an added hangnail can trigger strong *noninferiority*, even in turbulent lives. Consider again the case in which x_k ’s welfare is $\langle 9, 9 \rangle$ and x_{k+1} ’s is $\langle 9, 9.02 \rangle$. On $r_h = 0.5$, $w(x_k)_{r_h=0.5} = 0.5 \times 9 - 0.5 \times 9 = 0$ and $w(x_{k+1})_{r_h=0.5} = 0.5 \times 9 - 0.5 \times 9.02 = -0.01$. Adding zeroes can never yield a negative number, and vice versa, so x_k is strongly superior to x_{k+1} on $r_h = 0.5$: any number of lives x_k has greater value than any number of

²⁸ Reflecting this graph in the line $h = s$ gives the region of lives that can be pushed from bad to good by an increase of 0.02 in that life’s quantity of happiness. Perhaps this small jump corresponds to a gumdrop’s worth of pleasure. As above, the region includes only lives featuring very small quantities of happiness and suffering.

lives x_{k+1} on $r_h = 0.5$. That ensures that x_k is strongly noninferior to x_{k+1} overall: any number of lives x_k is not worse than any number of lives x_{k+1} .

More generally, an extra hangnail will trigger strong noninferiority whenever at least one of the lives being compared is weakly neutral. In that case, there will be some r_h on which the extra hangnail pushes the life's value from positive to negative. On that r_h , any number of lives without the hangnail has greater value than any number of lives with the hangnail. Therefore, any number of lives without the hangnail is not worse than any number of lives with the hangnail.

This too is a welcome result. Faced with a choice between adding two populations, at least one of which consists of lives that are neither good nor bad, it is not worse to choose the population containing the better lives, regardless of their respective sizes.

And, importantly, the IER-View does not imply strong noninferiority across straightforwardly-better-than-blank lives or strong nonsuperiority across straightforwardly-worse-than-blank lives, as Critical-Range Views do. To see why, consider a life y_k with welfare $\langle a, 0 \rangle$ and a life y_{k+1} with welfare $\langle a - 0.02, 0 \rangle$. Let $a > 0.02$, so that both y_k and y_{k+1} are straightforwardly-better-than-blank. In this case, both lives feature no suffering whatsoever, so $w(y_k)_r$ and $w(y_{k+1})_r$ are positive on each r in R . That implies that, for any number m , there is some number n such that n lives y_{k+1} have greater value than m lives y_k on each r . Therefore, n lives y_{k+1} are better than m lives y_k .

4. Less Concerning Greediness

Recall that Critical-Range Views imply Maximal Greediness: for each population of awful lives and each population of wonderful lives, (1) there is some population of straightforwardly-better-than-blank lives such that the population of awful lives is not worse than the population of wonderful lives plus the straightforwardly-better-than-blank lives, or (2) there is some population of straightforwardly-worse-than-blank lives such that the population of wonderful lives is not better than the population of awful lives plus the straightforwardly-worse-than-blank lives. This disjunction follows from Critical-Range Views' claim that lives at more than one welfare level are contributively weakly neutral and their assumption that any two lives are commensurable. Together, these imply that some straightforwardly-better-than-blank life or some straightforwardly-worse-than-blank life is contributively weakly neutral. And Critical-Range Views' commitment to Separability and Archimedeanism implies that adding enough contributively weakly neutral lives can make any population incommensurable with any other.

The IER-View agrees that lives at more than one welfare level are contributively weakly neutral. On the IER-View with $R = \{r: 0.4 \leq r_h \leq 0.6\}$, for example, lives at $\langle 4, 3 \rangle$ and $\langle 5, 4 \rangle$ are both weakly neutral. But, as we have

seen, it denies the assumption that any two lives are commensurable. $\langle 4, 3 \rangle$ and $\langle 5, 4 \rangle$ are one such incommensurable pair. As a result, it avoids Maximal Greediness. Blank lives – with welfare $\langle 0, 0 \rangle$ – have a value of 0 on each r in R , and so are contributively *strictly* neutral. Adding them to a population results in a population that is equally good, so they cannot swallow up goodness or badness.

Straightforwardly-better-than-blank lives, meanwhile – with welfare $\langle a, 0 \rangle$, $a > 0$ – have positive value on each r in R , and so are contributively good. Adding them to a population results in a population that is better, so they cannot swallow up and neutralise goodness. And *mutatis mutandis* for straightforwardly-worse-than-blank lives. They cannot swallow up and neutralise badness. Therefore, the IER-View implies neither disjunct of Maximal Greediness.

On the IER-View, only lives featuring some positive quantity of good can neutralise badness, and only lives featuring some positive quantity of bad can neutralise goodness. This is as it should be.

6. Objections to the Imprecise Exchange Rates View

The above four points constitute the main advantages of the IER-View. Below are two objections.

1. Some Incommensurability Between Good Lives and Weakly Neutral Lives

On the IER-View, some good lives are incommensurable with some weakly neutral lives. Take a life x with welfare $\langle 1, 0 \rangle$ and a life y with welfare $\langle 8, 7 \rangle$. Life x is good, because $w(x)_r$ is positive on each $0.4 \leq r_h \leq 0.6$. Life y is weakly neutral, because $w(y)_r$ is positive on each $r_h > 0.4\dot{6}$ and negative on each $r_h < 0.4\dot{6}$. Yet x is incommensurable with y , because $w(x)_r < w(y)_r$ on each $r_h > 0.5$ and $w(x)_r > w(y)_r$ on each $r_h < 0.5$.

Although this consequence might seem odd, we ought to accept it. The reasons are twofold. First, the implication is not unique to the IER-View but is an inevitable consequence of admitting the possibility of lives both weakly neutral and close-to-strictly neutral, as Gustafsson (2020, 96) and Rabinowicz (forthcoming, 86) note. To see how, recall that strictly neutral lives are equally good as the standard and weakly neutral lives are incommensurable with the standard. These definitions imply that strictly neutral lives are incommensurable with weakly neutral lives. As Raz (1986, 326) notes, a small improvement or detriment to either of two incommensurable objects typically does not remove their incommensurability. Such small tweaks can make a difference only when one of the two objects is almost better than the other. Therefore, if a strictly neutral life is neither almost better nor almost worse than some weakly neutral life, then

some good life – slightly better than the strictly neutral life – and some bad life – slightly worse than the strictly neutral life – will also be incommensurable with the weakly neutral life.

Second, this odd consequence follows from three claims that we should be reluctant to deny. The first is that a life featuring a positive quantity of good and no bad whatsoever – like $\langle 1, 0 \rangle$ – is good. The second is that a turbulent, neutral life – like $\langle 8, 7 \rangle$ – can be better than another neutral life – like $\langle 7, 7 \rangle$. The third is that the good life $\langle 1, 0 \rangle$ and the turbulent life $\langle 8, 7 \rangle$ are such that neither is better than the other and a small improvement either way fails to break the deadlock.

2. Some Instances of Maximal Repugnance

On the IER-View, life x with welfare $\langle a, 0 \rangle$ is good and life y with welfare $\langle 0, a \rangle$ is bad, for any $a > 0$. That implies that each population of wonderful lives is worse than some population of x -lives, and each population of awful lives is better than some population of y -lives. As a need only be larger than 0, lives x and y could be very similar. They could be identical but for x 's featuring an extra gumdrop and y 's featuring an extra hangnail. Therefore, the IER-View implies something like Maximal Repugnance. Gustafsson (2020, 96), Broome (forthcoming, 8), and Rabinowicz (forthcoming, 86–87) note that any view admitting the possibility of strictly neutral lives has implications of this kind, and they take it to be a reason to reject such views.

However, I claim that ruling out the IER-View on this basis is premature. Note first that implying this instance of Maximal Repugnance seems preferable to the alternative: claiming that lives with welfare $\langle a, 0 \rangle$ or $\langle 0, a \rangle$ for some $a > 0$ are contributively weakly neutral. As we have seen, that claim implies Maximal Greediness.

Note also that the IER-View implies Maximal Repugnance only when lives x and y are nearly-blank. If a life is turbulent, featuring a lot of happiness and suffering, then much more than a few extra gumdrops are required to move that life from bad to good. If we hold a life's quantity of suffering fixed at 6, for example, then the last contributively bad life has welfare $\langle 4, 6 \rangle$ and the first contributively good life has welfare $\langle 9, 6 \rangle$. Once again, the IER-View is giving gumdrops and hangnails their proper axiological due. In nearly-blank lives, they are significant. In turbulent lives, they fade into the background.

My final point is related. It is common in population axiology to imagine nearly-blank lives – like those at $\langle a, 0 \rangle$ – as drab. Parfit asked us to imagine lives in which the only pleasures are ‘muzak and potatoes’ (1986, 148). But we should note that if muzak-and-potatoes lives are to contain no bad whatsoever, then their protagonists must be very different from you and me. We – and everyone else endowed with an ordinary human psychology – would inevitably suffer

boredom and longing were we to live such a life. So, when we picture $\langle a, 0 \rangle$ lives, we should not imagine how we would feel sitting down to another bowl of mashed potatoes. Imagine instead a life of dreamless sleep, topped off with a gumdrop's worth of pleasure. When I conceive of $\langle a, 0 \rangle$ lives in this way, the IER-View's implications no longer strike me as so repugnant.

7. Conclusion

The variety of possible Critical-Set Views is dizzying, but each variety has serious drawbacks. On Critical-Level Views, just two hangnails can mark the difference between good and bad lives, even when the lives in question are long and turbulent. That means that a single life without the hangnails is better than any number of lives with them, and that each population of wonderful lives is worse than enough lives without the hangnails, while each population of awful lives is better than enough lives with them. On Critical-Range Views, meanwhile, each population of wonderful lives and each population of awful lives is such that adding enough lives featuring only good and neutral components to the former makes it no better than the latter, or adding enough lives featuring only bad and neutral components to the latter makes it no worse than the former. What's more, some discontinuity in contributive value must occur in a counterintuitive place, so that a single life featuring only dreamless sleep and some duration of happiness is not worse than any number of lives identical but for a slightly shorter duration of happiness, or a single life featuring only dreamless sleep and some duration of suffering is not better than any number of lives identical but for a slightly shorter duration of suffering. Some varieties of Critical-Level and Critical-Range Views are sadistic, and no variety can account for the incommensurability between lives and same-number populations without extra theoretical resources.

The Imprecise Exchange Rates View comes equipped with the required theoretical resources. It diagnoses as the source of incommensurability the fact that some trade-offs are neither worth it nor not worth it. The resulting incommensurability between lives allows us to claim both that blank lives are strictly neutral and that a wide range of turbulent lives are weakly neutral, so that the IER-View captures the respective advantages of both Critical-Level and Critical-Range Views and charts the narrow course between Maximal Greediness and the most concerning instances of Maximal Repugnance. Making the size of the contributively neutral range depend on a life's quantity of goods and bads has another nice consequence: it gives hangnails their proper axiological due. When a life is nearly-blank, an extra hangnail can take it from good to bad. When a life is turbulent, hangnails pale almost into axiological insignificance. And because the IER-View determines a life's personal and contributive value in the same way, it escapes all forms of sadism.

In sum, the IER-View is a worthy successor to Critical-Set Views. It retains much of their appeal, while avoiding many of their pitfalls.²⁹

8. References

- Arrhenius, Gustaf. 2000a. ‘An Impossibility Theorem for Welfarist Axiologies’. *Economics & Philosophy* 16 (2): 247–66.
- . 2000b. ‘Future Generations: A Challenge for Moral Theory’. PhD Thesis, Uppsala University.
- . 2011. ‘The Impossibility of a Satisfactory Population Ethics’. In *Descriptive and Normative Approaches to Human Behavior*, edited by Ehtibar N. Dzhafarov and Lacey Perry, 1–26. Singapore: World Scientific Publishing Company.
- Arrhenius, Gustaf, and Wlodek Rabinowicz. 2015a. ‘The Value of Existence’. In *The Oxford Handbook of Value Theory*, edited by Iwao Hirose and Jonas Olson, 424–44. New York: Oxford University Press.
- . 2015b. ‘Value Superiority’. In *The Oxford Handbook of Value Theory*, edited by Iwao Hirose and Jonas Olson, 225–48. New York: Oxford University Press.
- Blackorby, Charles, Walter Bossert, and David Donaldson. 1996. ‘Quasi-Orderings and Population Ethics’. *Social Choice and Welfare* 13 (2): 129–50.
- . 2005. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge: Cambridge University Press.
- Bossert, Walter. forthcoming. ‘Anonymous Welfarism, Critical-Level Principles, and the Repugnant and Sadistic Conclusions’. In *The Oxford Handbook of Population Ethics*, edited by Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns. Oxford: Oxford University Press.
- Broome, John. 1991. *Weighing Goods*. Oxford: Blackwell.
- . 2004. *Weighing Lives*. Oxford: Oxford University Press.
- . forthcoming. ‘Loosening the Betterness Ordering of Lives: A Response to Rabinowicz’. In *The Oxford Handbook of Population Ethics*, edited by Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns. Oxford: Oxford University Press. <http://users.ox.ac.uk/~sfop0060/pdf/Loosening%20the%20betterness%20ordering%20of%20lives.pdf>.

²⁹ I thank Hilary Greaves, Tomi Francis, Kacper Kowalczyk, Alice van’t Hoff, Todd Karhu, Nikhil Venkatesh, Jessica Fischer, Aidan Penn, Michal Masny, and two anonymous reviewers for helpful comments and discussion.

- Bykvist, Krister. 2007. 'The Good, the Bad and the Ethically Neutral'. *Economics & Philosophy* 23 (1): 97–105.
- Carlson, Erik. 1998. 'Mere Addition and Two Trilemmas of Population Ethics'. *Economics & Philosophy* 14 (2): 283–306.
- . forthcoming. 'On Some Impossibility Theorems in Population Ethics'. In *The Oxford Handbook of Population Ethics*, edited by Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns. Oxford: Oxford University Press.
- Chang, Ruth. 2002. 'The Possibility of Parity'. *Ethics* 112 (4): 659–88.
- De Sousa, Ronald B. 1974. 'The Good and the True'. *Mind* 83 (332): 534–51.
- Frick, Johann. 2017. 'On the Survival of Humanity'. *Canadian Journal of Philosophy* 47 (2–3): 344–67.
- Gustafsson, Johan E. 2020. 'Population Axiology and the Possibility of a Fourth Category of Absolute Value'. *Economics & Philosophy* 36 (1): 81–110.
- Hudson, James L. 1987. 'The Diminishing Marginal Value of Happy People'. *Philosophical Studies* 51 (1): 123–37.
- Huemer, Michael. 2008. 'In Defence of Repugnance'. *Mind* 117 (468): 899–933.
- Kitcher, Philip. 2000. 'Parfit's Puzzle'. *Noûs* 34 (4): 550–77.
- Meacham, Christopher J. G. 2012. 'Person-Affecting Views and Saturating Counterpart Relations'. *Philosophical Studies* 158 (2): 257–87.
- Nebel, Jacob M. 2021. 'Totalism without Repugnance'. In *Ethics and Existence: The Legacy of Derek Parfit*, edited by Jeff McMahan, Tim Campbell, James Goodrich, and Ketan Ramakrishnan. Oxford: Oxford University Press. <https://philpapers.org/archive/NEBTWR.pdf>.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- . 1986. 'Overpopulation and the Quality of Life'. In *Applied Ethics*, edited by Peter Singer, 145–64. Oxford: Oxford University Press.
- Qizilbash, Mozaffar. 2007. 'The Mere Addition Paradox, Parity and Vagueness'. *Philosophy and Phenomenological Research* 75 (1): 129–51.
- . 2018. 'On Parity and the Intuition of Neutrality'. *Economics & Philosophy* 34 (1): 87–108.
- Rabinowicz, Wlodek. 2009. 'Broome and the Intuition of Neutrality'. *Philosophical Issues* 19 (1): 389–411.
- . forthcoming. 'Getting Personal: The Intuition of Neutrality Reinterpreted'. In *The Oxford Handbook of Population Ethics*, edited by Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns. Oxford: Oxford University Press. <https://www.iffs.se/en/publications/working-papers/studies-on-climate-ethics-and-future-generations-vol-2/>.
- Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Oxford University Press.

- Tännsjö, Torbjörn. 2002. 'Why We Ought to Accept the Repugnant Conclusion'. *Utilitas* 14 (3): 339–59.
- Thomas, Teruji. 2018. 'Some Possibilities in Population Axiology'. *Mind* 127 (507): 807–32.
- . forthcoming. 'Separability and Population Ethics'. In *The Oxford Handbook of Population Ethics*, edited by Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns. Oxford: Oxford University Press.
- Thornley, Elliott. forthcoming. 'A Dilemma for Lexical and Archimedean Views in Population Axiology'. *Economics & Philosophy*.
- . 2021. 'The Impossibility of a Satisfactory Population Prospect Axiology (Independently of Finite Fine-Grainedness)'. *Philosophical Studies*. <https://doi.org/10.1007/s11098-021-01621-4>