

CRITICAL LEVELS, CRITICAL RANGES, AND IMPRECISE EXCHANGE RATES IN POPULATION AXIOLOGY

Elliott Thornley

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 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 subtract a range of positive constants from the number representing a person's welfare and then calculate the value of a population relative to each constant within the range. 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 constant within the 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 characterization 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 at least one implausible discontinuity, and no view can account for the incommensurability between lives and between 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 aforementioned 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 between same-size populations, avoids all forms of sadism along with the most concerning instances of repugnance and greediness, and has many other advantages besides.

I characterize and taxonomize critical-set views in section 1 and object to them in section 2. I introduce the Imprecise Exchange Rates View in section 3, canvas its advantages in section 4, and address some objections in section 5. I sum up in section 6.

1. CRITICAL-SET VIEWS

Foundational to critical-set views is the notion of a *life*. I follow Broome 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.¹ 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 both measurable on an interval scale and interpersonally level 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 the life Ada would have as an artist and the life Ada would have as a baker is twice the size of the difference in welfare between the life Ada would have as a baker

1 Broome, *Weighing Lives*, 94–95.

and the life Ada would have as a consultant.” Interpersonal level comparability allows us to compare the welfare of different people, so that claims like the following are meaningful: “The life Ada would have as an artist contains more welfare than the life Bob would have as a baker.” This claim is equivalent to the claim that “The life Ada would have as an artist is personally better than the life Bob would have as a baker.” In other words, “The life Ada would have as an artist is better *for her* than the life Bob would have as a baker is *for him*.” I mostly use the terminology of personal betterness below.

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 $|w(x) - w(y)| = k|w(y) - w(z)|$. 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 . I will call $w(x)$ the *welfare level* of life 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. Following Rabinowicz, I will say that 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.² The standard in question is defined differently by different authors. Some define it as nonexistence.³ Others define it as a life constantly at a neutral level of temporal welfare.⁴ 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.⁵ With one caveat, critical-set views are compatible with each definition.⁶

So much for comparing lives. Comparing populations—sets of lives—requires more machinery. Critical-set views start by designating some (gapless)

- 2 Rabinowicz, “Getting Personal,” 80–81. Gustafsson calls these lives “neutral” and “undistinguished,” respectively (“Population Axiology and the Possibility of a Fourth Category of Absolute Value”).
- 3 Arrhenius and Rabinowicz, “The Value of Existence.”
- 4 Broome, *Weighing Lives*, 68; Bykvist, “The Good, the Bad, and the Ethically Neutral,” 101.
- 5 Arrhenius, “Future Generations,” 26.
- 6 The caveat is that *neutral-range views*—explained below—cannot be paired with the latter two definitions. Neutral-range views claim 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.

set of welfare levels to be the *critical set*. This critical set is defined to be the set of all welfare levels such that adding lives at these welfare levels to a population makes that population neither better nor worse. Each welfare level within this critical set is called a *critical level*. These critical levels play a key role in determining a life's *contributive value*, which we can understand as the contribution that a life makes to the value of a population. On critical-set views, the contributive value $c(x)_q$ of a life x relative to a critical level q is calculated by subtracting q from the welfare level $w(x)$:⁷

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

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 $v(X)_q \geq v(Y)_q$ 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.

Here is an example to illustrate. Suppose that we have two populations, X and Y . X contains one person at welfare level 5. Y contains three people at welfare level 2. 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 characterization prior to this example constitutes the common core of critical-set views. The following four 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. The former are *critical-level views*, and the latter are *critical-range views*. On critical-level views, lives at the critical level are *contributively strictly neutral*, by which I mean that adding these lives to a population leaves the new population equally good as the original. On critical-range

7 Critical-set views can also incorporate some real-valued function f applied to the welfare level and critical level. This function could be prioritarian: strictly increasing and strictly concave. I leave out the f purely for simplicity's sake. My discussion applies to any critical-set view on which f is strictly increasing. Any critical-set view on which f is not strictly increasing will violate *Dominance over Persons*, which says that for any populations X and Y featuring all the same people, if each person is at least as well off in X as they are in Y and some person is better off in X than they are in Y , then X is better than Y .

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 $v(Y)_4 = (2 - 4) + (2 - 4) + (2 - 4) = -6$.

views, lives within the critical range are *contributively weakly neutral*, by which I mean that adding these lives to a population renders the new population incommensurable with the original. On all critical-set views, adding lives at welfare levels above the critical set makes a population better and adding lives at welfare levels below the critical set makes a population worse. I will call such lives *contributively good* and *contributively bad*, respectively.

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 personally neutral level are personally *strictly* neutral, so that they are personally equally good as the standard. *Neutral-range views* claim that lives within the personally neutral range are personally *weakly* neutral, so that they are personally incommensurable with the standard. From now on, I drop the “personally” from expressions like “personally neutral set.” “Neutral set” refers to the set of welfare levels such that lives at those levels are personally neutral. “Critical set” refers to the set of welfare levels such that lives at those levels 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, parity, indeterminacy, some other value relation, or any combination of the aforementioned phenomena.¹⁰ 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 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.

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 lives at the corresponding welfare level are personally/contributively *strictly* neutral. Boxes denote that lives at the corresponding welfare levels are personally/contributively *weakly* neutral. Lives at welfare levels above (below) the horizontal line or shaded box are personally/contributively good (bad). The numbers are purely illustrative.

10 For incommensurability, see Blackorby, Bossert, and Donaldson, “Quasi-Orderings and Population Ethics.” For parity, see Qizilbash, “The Mere Addition Paradox, Parity and Vagueness” and “On Parity and the Intuition of Neutrality”; and Rabinowicz, “Broome and the Intuition of Neutrality.” For indeterminacy, see Broome, *Weighing Lives*.

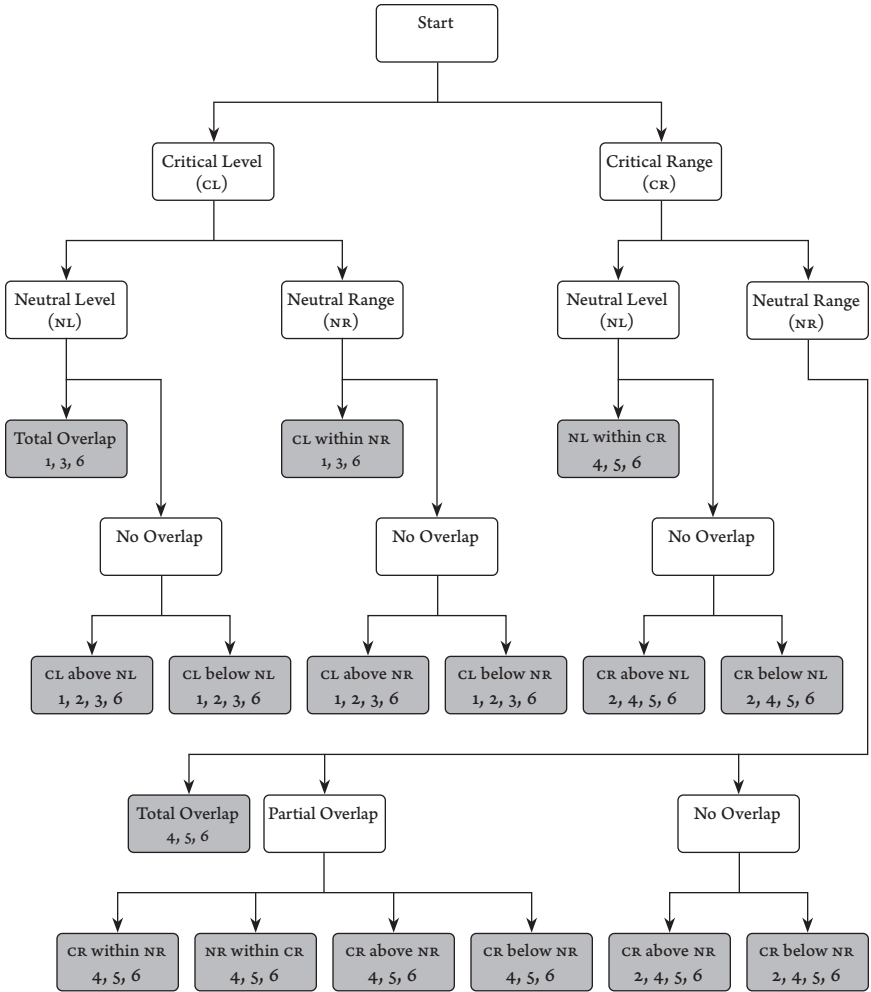


FIGURE 1 Taxonomy

First, the Total View (fig. 2), which is defended by Hudson, Tännsjö, and Huemer, among others.¹¹ There is a single coinciding 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.

11 Hudson, "The Diminishing Marginal Value of Happy People"; Tännsjö, "Why We Ought to Accept the Repugnant Conclusion"; Huemer, "In Defence of Repugnance."




FIGURE 2 The total view

Second, a positive critical-level view (fig. 3), defended by Blackorby, Bossert, and Donaldson.¹² There is a single critical level above a single neutral level, so a life can be personally good without being contributively good. Any two populations are commensurable.




FIGURE 3 A positive critical-level view

12 Blackorby, Bossert, and Donaldson, *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*; Bossert, "Anonymous Welfarism, Critical-Level Principles, and the Repugnant and Sadistic Conclusions."

Third, a critical-range view. A view of this kind is defended by Broome, who interprets the critical range as a range of indeterminacy, along with Qizilbash and Rabinowicz, who each interpret the critical range as a range of parity.¹³ There is a single neutral level but a critical range, so any overlap between the neutral and critical sets can be partial at most. In figure 4, I present a version of the view in which the neutral level coincides with the lowest welfare level in the critical range. On critical-range views, some pairs of populations are incommensurable.

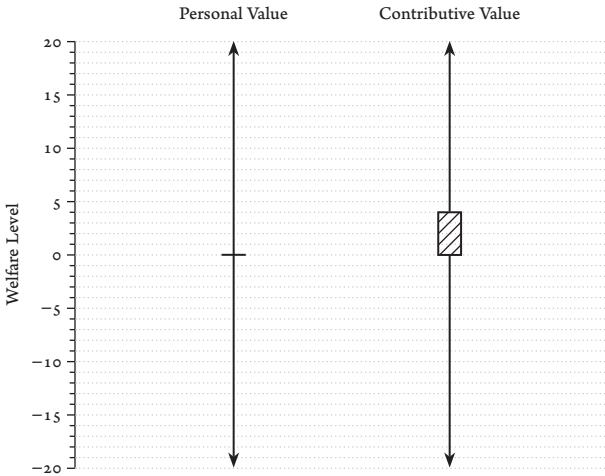


FIGURE 4 A critical-range view

Finally, a neutral-range view (fig. 5). Rabinowicz discusses a view of this kind in more recent work, and Gustafsson 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.¹⁴ 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.

13 Broome, *Weighing Lives*; Qizilbash, “The Mere Addition Paradox, Parity and Vagueness” and “On Parity and the Intuition of Neutrality”; Rabinowicz, “Broome and the Intuition of Neutrality.”
 14 Rabinowicz, “Getting Personal”; Gustafsson, “Population Axiology and the Possibility of a Fourth Category of Absolute Value.”




FIGURE 5 A neutral-range view

2. OBJECTIONS TO CRITICAL-SET VIEWS

Many varieties of critical-set view are subject to the same objections. Each view must reckon with at least three of the following six.

2.1. *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.¹⁵

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.¹⁶

Both of these consequences arise because, on critical-set views, a population of

15 Parfit, *Reasons and Persons*, 388.

16 Gustafsson, "Population Axiology and the Possibility of a Fourth Category of Absolute Value," 85. Carlson calls this claim the "Reverse Repugnant Conclusion" ("Mere Addition and Two Trilemmas of Population Ethics," 297). Broome calls it the "Negative Repugnant Conclusion" (*Weighing Lives*, 213).

enough contributively good (bad) lives can be better (worse) than any other population.

However, as Rabinowicz notes, the repugnance of these conclusions is attenuated if lives at a wide range of welfare levels are personally neutral.¹⁷ 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:

Maximal Repugnance: There is a life x and a life y that is identical but for one fewer gumdrop's worth of pleasure and one more 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 that 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 one fewer gumdrop's worth of pleasure and one extra hangnail's worth of pain can push a life's welfare level from above the critical level to below it, all critical-level views imply Maximal Repugnance.

2.2. Sadism

Any view on which there is no overlap between the critical set and the neutral set implies some sadistic conclusion. If the critical set is above the neutral set and there is some welfare level between the two, the view implies the original

Sadistic Conclusion: Each population of awful lives is better than some population of personally good lives.¹⁸

That is because lives at a welfare level above the neutral set and below the critical

17 Rabinowicz, "Broome and the Intuition of Neutrality," 406, and "Getting Personal," 79.

18 Arrhenius, "An Impossibility Theorem for Welfarist Axiologies," 256.

set are personally good but contributively bad. And on critical-set views, adding enough contributively bad lives to a population can make that population worse than any other.

If the critical set is below the neutral set and there is some welfare level between them, the view implies the

Mirrored Sadistic Conclusion: Each population of wonderful lives is worse than some population of personally bad lives.¹⁹

That is because lives at a welfare level below the neutral set and above the critical set are personally bad but contributively good. And on critical-set views, adding enough contributively good lives to a population can make that population better than any other.

We could endorse a critical-set view on which there is no overlap between the neutral set and the critical set and yet no welfare level between the two sets.²⁰ These kinds of views imply only weaker forms of sadism. If the critical set is above 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 critical set is below 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.²¹

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

- 19 Gustafsson, "Population Axiology and the Possibility of a Fourth Category of Absolute Value," 85.
- 20 That is possible if welfare levels are *not dense* (by which I mean that there is some pair of distinct welfare levels with no welfare level between them) or if the neutral set and critical set are such that exactly one of them is open at the end where they meet (e.g., if the neutral set is $[0, 1)$ and the critical set is $[1, 2]$).
- 21 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 ("Population Axiology and the Possibility of a Fourth Category of Absolute Value," 86) and Meacham ("Person-Affecting Views and Saturating Counterpart Relations," 270), respectively.

2.3. 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 extra hangnail's worth of pain. We reach x_3 by making x_2 slightly worse, and so on. 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 nonpositive. Adding positive numbers can never yield a nonpositive number, and vice versa, so critical-level views imply that any population of lives x_k is better than any population of lives x_{k+1} . Call this implication *Strong Superiority across Slight Differences* (SSASD).²²

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. Consider a case in which each life in our x -sequence is long and turbulent, featuring soaring highs and crushing lows. Amid 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 can make all the difference: there will be a long, turbulent life x_k such that any population of lives x_k is better than any population of lives x_{k+1} identical but for the extra hangnail. Two corollaries of this implication bring out its implausibility: a population of just a single life without the hangnail is better than any population of lives with it, and a population of just a single life with the hangnail is worse than any population of lives without it.

2.4. Strong Noninferiority across Slight Differences

This instance of SSASD might spur us to adopt a critical-range view. On critical-range views, lives at a range of welfare levels 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 population of contributively good lives is better than any population of contributively weakly neutral lives. Instead, each population of con-

22 For discussions of superiority and noninferiority in axiology, see Arrhenius and Rabinowicz, "Value Superiority"; Nebel, "Totalism without Repugnance"; and Thornley, "A Dilemma for Lexical and Archimedean Views in Population Axiology."

tributively good lives is incommensurable with some population of contributively weakly neutral lives. Here is an example to warm us up for the proof.

Suppose that all the welfare levels between 0 and 4 inclusive are critical. 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 critical level q in the critical set Q . But X is incommensurable with population Z consisting of two lives x_{k+1} . X has greater value than Z relative to $q = 4$, but Z has greater value than X relative to $q = 0$.²³

More generally, each contributively weakly neutral life has positive contributive value relative to some critical level q .²⁴ That implies that each population has less value than some sufficiently large population of contributively weakly neutral lives relative to that q . Therefore, each population is not better than some sufficiently large population of 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 population of lives x_k is *not worse than* any population of lives x_{k+1} . To see how, return to our example above. 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 $q = 4$. Therefore X is not worse than Z , no matter what their respective sizes. More generally, for any contributively good life x_k and any contributively weakly neutral life x_{k+1} , there exists some q such that x_k has positive contributive value relative to q and x_{k+1} has nonpositive contributive value relative to q . So relative to this q , any population of lives x_k has greater value than any population of lives x_{k+1} . That in turn implies that any population of lives x_k is not worse than any population of lives x_{k+1} . This kind of discontinuity is innocuous considered in itself. But as I demonstrate below, critical-range views imply that Strong Noninferiority across Slight Differences occurs in some counterintuitive places.

23 $v(X)_4 = (4.01 - 4) = 0.01$ and $v(Z)_4 = (3.99 - 4) + (3.99 - 4) = -0.02$; $v(X)_0 = (4.01 - 0) = 4.01$ and $v(Z)_0 = (3.99 - 0) + (3.99 - 0) = 7.98$.

24 We might think that lives at the lowest welfare level in the critical range are a counterexample to this claim. They do not have positive value relative to any critical level q in the critical range Q . But these lives are not contributively weakly neutral. On our definitions, they are contributively bad. Here is why. Suppose $w(x)$ is the lowest welfare level in the critical range Q . Then, for any population X , the value of X is at least as great as the value of X plus a life at $w(x)$ relative to each q in Q , so X is at least as good as X plus a life at $w(x)$. But the value of X plus a life at $w(x)$ is *not* at least as great as the value of X relative to each q in Q (in particular, it is not at least as great relative to critical levels q that are not the lowest in the critical range), so X plus a life at $w(x)$ is not at least as good as X . Therefore, X plus a life at $w(x)$ is worse than X , 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.

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*. Life y_n features only the blank period and so qualifies as a *blank life*, featuring no good or bad components whatsoever.²⁵ Life 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 population of lives y_{n-1} featuring one hour of happiness is not worse than any population 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 population of lives y_{n+1} featuring one hour of suffering is *not better* than any population 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 critical. Therefore, in any sequence with sufficiently small differences in welfare between adjacent lives, more than one life is contributively weakly neutral. We can make the differences in welfare 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, any population of lives featuring two seconds of happiness is not worse than any population of lives featuring one second of happiness. That implies that a population of just a *single* life featuring two seconds of happiness is not worse than any population of lives featuring one second of happiness. But this consequence seems implausible. The only difference between the lives is the duration of happiness; the latter population can feature an arbitrarily longer total duration of happiness, and yet the latter population can never be better than the former.

25 Broome, *Weighing Lives*, 208.

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

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

2.5. *Maximal Greediness*

Critical-range views face another difficulty. As Broome points out, they imply that contributively weakly neutral lives can “swallow up” and neutralize goodness and badness.²⁶ Here is an illustration of what that means. Suppose again that all welfare levels between 0 and 4 inclusive are critical. 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 level 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 weakly neutral. But our critical-range view implies that these changes are not bad. Neither *A*’s nor *B*’s value is at least as great as the other relative to each q in Q , so the two populations are incommensurable.²⁷ Our critical-range view also implies that *A* is incommensurable with *C* (in which *x*’s welfare level is 18 and there are two lives at welfare level 2) and *D* (in which *x*’s welfare level 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 level is extremely low and there is some large number of contributively weakly neutral lives. Broome and I find this “greedy neutrality” concerning, but others are happy to bite the bullet.²⁸ In any case, the worry can be sharpened.

Note first that the size of population *A* need not be restricted to a single life:

26 Broome, *Weighing Lives*, 169–70 and 202–5.

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

28 Rabinowicz, “Broome and the Intuition of Neutrality”; Frick, “On the Survival of Humanity”; Gustafsson, “Population Axiology and the Possibility of a Fourth Category of Absolute Value.”

adding enough contributively weakly neutral lives can neutralize 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 population. The former population might feature 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*.

2.6. *No Incommensurability between Lives or between Same-Size 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 fol-

lowing small improvement argument.²⁹ Suppose that X consists of ten wonderful lives and Y consists of one hundred very good lives. Neither X nor Y is better than the other.³⁰ 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 one hundred slightly-better-than-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-size 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 .

Now note critical-set views' equation for the value of a population X relative to a critical level q :

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

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

1. For any populations X and Y and any critical 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-size populations and q stand for an arbitrary critical 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:

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

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

$$(\sum_i w(x_i)) - nq \geq (\sum_i w(y_i)) - nq.$$

29 De Sousa, "The Good and the True"; Chang, "The Possibility of Parity."

30 Those who disagree should tweak the numbers or adjectives.

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

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

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-size populations X and Y and any critical 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-size populations X and Y , X is at least as good as Y or Y is at least as good as X .

In other words, critical-set views imply that any two same-size populations are commensurable.

However, universal commensurability of same-size 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.³¹ 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 suggest the incommensurability of other pairs of same-size populations.

Partly on the basis of such arguments, advocates of critical-set views have started to incorporate incommensurability and indeterminacy into their theories of personal betterness. Broome, for example, states that some pairs of lives are obviously indeterminately related but offers no explanation for why this is so.³² Rabinowicz, 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.³³ And Gustafsson accounts for incommensurability between lives by claiming that there is a neutral range of temporal welfare levels.³⁴ 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

31 Parfit, "Overpopulation and the Quality of Life," 148.

32 Broome, "Loosening the Betterness Ordering of Lives."

33 Rabinowicz, "Getting Personal."

34 Gustafsson, "Population Axiology and the Possibility of a Fourth Category of Absolute Value."

the same reason that critical-range views cannot account for the incommensurability between same-size 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 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.³⁵ This notion of "different weight assignments" forms the core of the Imprecise Exchange Rates View.

3. IMPRECISE EXCHANGE RATES

Some trade-offs are worth making. For example, going to the dentist to prevent tooth decay is a trade-off worth making. The good of having healthy teeth outweighs the bad of the trip. Other trade-offs are worth *not* making. Getting up at 4 AM and walking to work to save the £2 bus fare is a trade-off worth not making. The bad outweighs the good. Still other trade-offs are neither worth making nor worth not making, and a small improvement fails to break the deadlock. Here is an example.

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 peas and can see one hundred 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, the trade-off may be neither worth making nor worth not making. 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 pea from the plate. That need not ensure that finishing dinner is now a trade-off worth making.

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 peas and the jelly is incommensurable with eating neither. This incommensurability between goods, bads, and their combinations

35 Rabinowicz, "Getting Personal," 81.

is the source of incommensurability between lives. The child's life in which they eat the peas and jelly is incommensurable with the otherwise identical life in which they eat neither.

That is one motivation for the Imprecise Exchange Rates (IER) view. Now for the formalization. Recall that critical-set views begin with an ordering of lives by welfare. 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 (h), love (l), suffering (s), and false belief (f) are the dimensions of good and bad. Then the welfare level of a life x is as follows:

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

I assume that h , l , s , and f are real-valued functions. I also assume that the values of each function are interpersonally level comparable (so that we can make claims like "The life Ada would have as an artist features more happiness than the life Bob would have as a baker") and measurable on a ratio scale (so that we can make claims like "The life Ada would have as an artist features twice the suffering of the life Ada would have 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 the life Ada would have 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 the life Ada would have as an artist, the claim that her happiness outweighs her suffering relative to a given r will be true iff $r_h h(x) > r_s s(x)$.

On the IER view, only welfare levels *relative to a given* r can be expressed 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 relative to r is the sum of the welfare levels of each of its lives relative to 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 admissible proto-exchange rates R . A life x is at least as good as a life y iff $w(x)_r \geq w(y)_r$ relative to each r in R . And a population X is at least as good as a population Y iff $w(X)_r \geq w(Y)_r$ relative to each r in R .³⁶

In what follows, I mostly discuss a simple hedonist version of the IER view, in which the welfare level of a life x is given by a vector of happiness and suffering, $\langle h(x), s(x) \rangle$, with the functions h and s normalized so that the proto-exchange-rate r composed of $r_h = 0.5$ and $r_s = 0.5$ falls within the set R . 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.

4. ADVANTAGES OF THE IMPRECISE EXCHANGE RATES VIEW

The IER view has several advantages over critical-set views. Here are four.

4.1. *Some Incommensurability between Lives and between Same-Size Populations*

The first advantage is that the IER view offers a simple and plausible account of incommensurability between lives and between same-size populations. Recall that a life is at least as good as another iff its welfare level is at least as great relative to 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 welfare level of x is greater relative to $r_h = 0.4$, but the welfare level of y is greater relative to $r_h = 0.6$.³⁷ This is as it should be. Taking on the extra suffering in y for the sake of the extra happiness is a trade-off neither worth making nor worth not making.

The IER view also gives us the right result in small improvement cases. A

36 Rabinowicz offers a similar formalization (“Getting Personal,” 83–84). His formalization, 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.

37 $w(x)_{r_h=0.4} = 0.4 \times 4 - 0.6 \times 1 = 1$ and $w(y)_{r_h=0.4} = 0.4 \times 10 - 0.6 \times 6 = 0.4$; $w(x)_{r_h=0.6} = 0.6 \times 4 - 0.4 \times 1 = 2$ and $w(y)_{r_h=0.6} = 0.6 \times 10 - 0.4 \times 6 = 3.6$.

slightly improved life y^+ at welfare level $\langle 10 + \epsilon, 6 \rangle$ comes out better than y and incommensurable with x . That is because the IER view accounts for the 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 .³⁸

Another implication is related. Let us say that two proto-exchange rates *differ in optimism* iff they differ in the total weight granted to all dimensions of good taken together.³⁹ The implication is that if R contains proto-exchange rates that differ in optimism, then only lives featuring identical quantities of good and bad can be equally good.⁴⁰ That means that lives at welfare levels such as $\langle 4, 4 \rangle$ and

38 Here is a sketch of the proof. Life x is at least as good as life y relative to 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 nonnegative. 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.

39 Here is an example. Return briefly to our objective list theory on 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 $r_3 = \langle 0.3, 0.3, 0.1, 0.3 \rangle$. Proto-exchange-rates 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 of good taken together. Proto-exchange-rate 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.

40 To see this result, note first that equally good lives must have the same welfare level relative to each proto-exchange-rate. If x has a greater welfare level than y relative to 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 in x , and so on, and let r_1 and r_2 denote the total weight assigned to dimensions of good relative to 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. Canceling $b(x) - b(y)$ from each side gives

$\langle 5, 5 \rangle$ 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 a trade-off neither worth making nor worth not making. If lives at $\langle 4, 4 \rangle$ and $\langle 5, 5 \rangle$ were judged equally good, the view would generate counterintuitive verdicts in small improvement cases. For example, a life at $\langle 4, 4 \rangle$ would be worse than a life at $\langle 5, 5 - \epsilon \rangle$ for any $\epsilon > 0$. From now on, 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-size 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 quantities of good and bad can be equally good.

4.2. 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 or bad can be equally good. The extension of personal strict neutrality 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 or bad in the population, and no populations differing in their quantities of good or bad can be equally good.

This coincidence of personal and contributive strict neutrality suffices to establish that each category of personal value coincides with the corresponding category of contributive value. That is because the IER view then determines

$$r_1(g(x) - g(y) + b(x) - b(y)) = r_2(g(x) - g(y) + b(x) - b(y)).$$

Since $r_1 \neq r_2$, the expression $g(x) - g(y) + b(x) - b(y)$ must equal 0. That is true iff there exists some k such that $g(x) - g(y) = k$ 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.

each life’s personal and contributive category in the same way: its value is compared to the value of a blank life relative to 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 each personal and contributive category of value on the IER view established, I often drop the words “personal” and “contributive” in what follows. In figure 6, I graph these coincident categories for lives at different welfare levels 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 dark (light/white) region. Lives at the origin are blank and hence strictly neutral.

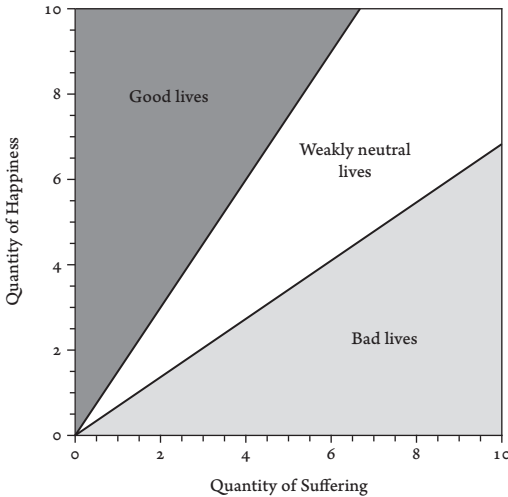


FIGURE 6 Coincident categories at different welfare levels.

4.3. Less Concerning Superiority and Noninferiority

As we saw above, critical-level views imply a concerning instance of Strong Superiority across Slight Differences (SSASD) in our x -sequence: there exists some long, turbulent life x_k such that any population of lives x_k is better than any population of lives x_{k+1} identical but for an extra hangnail. Critical-range views, meanwhile, imply only Strong Noninferiority across Slight Differences in our x -sequence: there exists some long, turbulent life x_k such that any population of lives x_k is not worse than any population 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 there exists some

life y_k featuring only neutral components and happiness such that a population of just a single life y_k is not worse than any population of lives each featuring a slightly shorter duration of happiness, or there exists some life y_j featuring only neutral components and suffering such that a population of just a single life y_j is not better than any population 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 some turbulent life x_k has welfare level $\langle 9, 9 \rangle$. Life x_{k+1} then has welfare level $\langle 9, 9.02 \rangle$. 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 population Z , consisting of two lives x_{k+1} . X has greater value than Z relative to $r_h = 0.4$, but Z has greater value than X relative to $r_h = 0.6$.⁴¹

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 $\langle h(x_k), s(x_k) + 0.02 \rangle$ be the welfare levels of x_k and x_{k+1} respectively. Life x_k is strongly superior to life 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. This condition is satisfied iff x_k 's welfare level is nonnegative relative to the most pessimistic proto-exchange rate $r_h = 0.4$, x_{k+1} 's welfare level is nonpositive relative to the most optimistic proto-exchange rate $r_h = 0.6$, and at least one of x_k 's or x_{k+1} 's welfare levels is non-zero relative to some r in R .⁴² 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 gives us the region in 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 can be consequential. In turbulent lives, they pale almost into axiological insignificance.⁴³

41 $v(X)_{r_h=0.4} = 0.4 \times 9 - 0.6 \times 9 = -1.8$ and $v(Z)_{r_h=0.4} = (0.4 \times 9 - 0.6 \times 9.02) + (0.4 \times 9 - 0.6 \times 9.02) = -3.624$; $v(X)_{r_h=0.6} = 0.6 \times 9 - 0.4 \times 9 = 1.8$ and $v(Z)_{r_h=0.6} = (0.6 \times 9 - 0.4 \times 9.02) + (0.6 \times 9 - 0.4 \times 9.02) = 3.584$.

42 The hangnail's worth of pain ensures that this last condition is met.

43 Reflecting this graph in the line $h = s$ gives the region of lives that can be pushed from bad or strictly neutral 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 in figure 7, the region includes only lives featuring very small quantities of happiness and suffering.




FIGURE 7 Welfare levels at which an extra hangnail triggers strong superiority

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 level is $\langle 9, 9 \rangle$ and x_{k+1} 's welfare level is $\langle 9, 9.02 \rangle$. Given $r_h = 0.5$, $w(x_k)_{r_h=0.5} = 0.5 \times 9 - 0.5 \times 9 = 0$, 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 any population of lives x_k has greater value than any population of lives x_{k+1} relative to $r_h = 0.5$. That ensures that x_k is strongly noninferior to x_{k+1} : any population of lives x_k is not worse than any population 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, the extra hangnail will push the life's value from positive to negative relative to some r_h . Relative to that r_h , any population of lives without the hangnail has greater value than any population of lives with the hangnail. Therefore, any population of lives without the hangnail is not worse than any population of lives with the hangnail.

This too is a welcome result. Suppose we must choose between two populations. Each population consists of lives at only one welfare level, one population's lives are better than the other's, and at least one population consists of lives that are neither good nor bad. Then it is not worse to choose the population consisting of the better lives, regardless of the populations' 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 level $\langle a, 0 \rangle$ and a life y_{k+1} with welfare level $\langle b, 0 \rangle$. Suppose that $a > b > 0$, so that y_k is better than y_{k+1} and both are straightforwardly better than blank. Since both lives feature no suffering whatsoever, $w(y_k)_r$ and $w(y_{k+1})_r$ are positive relative to each r in R . That implies that for any r in R and any number m , there is some number n such that a population of n lives y_{k+1} has greater value than a population of m lives y_k relative to r . So for any number m , there is some number n such that a population of n lives y_{k+1} is better than a population of m lives y_k . The result is that y_k is not strongly noninferior to y_{k+1} .⁴⁴ A parallel line of argument proves that no straightforwardly worse-than-blank life is strongly nonsuperior to any other straightforwardly worse-than-blank life.

4.4. Less Concerning Greediness

Recall that critical-range views imply Maximal Greediness: for any population of awful lives and any 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 on critical-range views, adding enough contributively weakly neutral lives to a population can make that 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. Lives at $\langle 4, 3 \rangle$ and $\langle 5, 4 \rangle$ are one such incommensurable pair. As a result, the IER view avoids Maximal Greediness. Blank lives—with welfare level $\langle 0, 0 \rangle$ —have a value of 0 relative to each r in R , and so are contributively *strictly* neutral. Adding them to a population leaves the new population equally good as the original, so blank lives cannot swallow up goodness or badness.

Straightforwardly better-than-blank lives, meanwhile—with welfare level

44 Indeed, y_k is not even *weakly* noninferior to y_{k+1} . For the distinction between strong and weak noninferiority, see Thornley, "A Dilemma for Lexical and Archimedean Views in Population Axiology," 6.

$\langle a, 0 \rangle$, $a > 0$ —have positive value relative to each r in R , and so are contributively good. Adding them improves a population, so straightforwardly better-than-blank lives cannot swallow up and neutralize goodness. And *mutatis mutandis* for straightforwardly worse-than-blank lives. They cannot swallow up and neutralize badness. Therefore, the IER view implies neither disjunct of Maximal Greediness.

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

5. OBJECTIONS TO THE IMPRECISE EXCHANGE RATES VIEW

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

5.1. 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 level $\langle 1, 0 \rangle$ and a life y with welfare level $\langle 8, 7 \rangle$. Life x is good, because $w(x)_r$ is positive relative to each $0.4 \leq r_h \leq 0.6$. Life y is weakly neutral, because $w(y)_r$ is positive relative to each $r_h > 0.4\dot{6}$ and negative relative to each $r_h < 0.4\dot{6}$. Yet x is incommensurable with y , because $w(x)_r < w(y)_r$ relative to each $r_h > 0.5$ and $w(x)_r > w(y)_r$ relative to 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. It is an inevitable consequence of admitting the possibility of lives both weakly neutral and close-to-strictly neutral, as Gustafsson and Rabinowicz note.⁴⁵ To see why, recall that strictly neutral lives are equally good as the standard and that weakly neutral lives are incommensurable with the standard. These definitions imply that strictly neutral lives are incommensurable with weakly neutral lives. As Raz notes, a small improvement or detriment to either of two incommensurable objects typically does not remove their incommensurability.⁴⁶ 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.

45 Gustafsson, "Population Axiology and the Possibility of a Fourth Category of Absolute Value," 96; Rabinowicz, "Getting Personal," 86.

46 Raz, *The Morality of Freedom*, 326.

Second, incommensurability between some good lives and some weakly neutral lives 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 a life at welfare level $\langle 1, 0 \rangle$) is good. The second is that a turbulent, neutral life (like a life at welfare level $\langle 8, 7 \rangle$) can be better than another neutral life (like a life at welfare level $\langle 7, 7 \rangle$). The third is that a good life at welfare level $\langle 1, 0 \rangle$ and a turbulent life at welfare level $\langle 8, 7 \rangle$ are such that neither is better than the other and a small improvement either way fails to break the deadlock.

5.2. *Some Instances of Maximal Repugnance*

On the IER view, life x with welfare level $\langle a, 0 \rangle$ is good and life y with welfare level $\langle 0, a \rangle$ 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 Maximal Repugnance. Gustafsson, Broome, and Rabinowicz 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.⁴⁷

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, which is to claim that lives with welfare level $\langle a, 0 \rangle$ or $\langle 0, a \rangle$ for some $a > 0$ are contributively weakly neutral. As we have seen, that claim commits critical-set views to 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 level $\langle 4, 6 \rangle$ and the first contributively good life has welfare level $\langle 9, 6 \rangle$. 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 think of lives barely worth living as drab. Parfit asked us to imagine lives in which the only pleasures are "muzak and potatoes."⁴⁸ But a Muzak and potatoes life can have

47 Gustafsson, "Population Axiology and the Possibility of a Fourth Category of Absolute Value," 96; Broome, "Loosening the Betterness Ordering of Lives," 8; Rabinowicz, "Getting Personal," 86–87.

48 Parfit, "Overpopulation and the Quality of Life," 148.

a welfare level of $\langle a, o \rangle$ only if its protagonist is very different from you and me. We—and everyone else endowed with an ordinary human psychology—would inevitably suffer boredom were we to live such a life, and lives at welfare level $\langle a, o \rangle$ feature no bad whatsoever. So, when we picture lives at $\langle a, o \rangle$, 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 $\langle a, o \rangle$ lives in this way, the IER view's implications no longer strike me as so repugnant.

6. CONCLUSION

The variety of possible critical-set views is dizzying, but each variety has serious drawbacks. On critical-level views, two extra hangnails can mark the difference between a good life and a bad life, even when the lives in question are long and turbulent. That means that a population of just a single life without the hangnails is better than any population of lives with them. It also means that each population of wonderful lives is worse than some population of lives without the hangnails, while each population of awful lives is better than some population of 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 is more, some discontinuity in contributive value must occur in a counterintuitive place, so that a population of just a single life featuring only dreamless sleep and some duration of happiness is not worse than any population of lives identical but for a slightly shorter duration of happiness, or a population of just a single life featuring only dreamless sleep and some duration of suffering is not better than any population of lives identical but for a slightly shorter duration of suffering. Some varieties of critical-set views are sadistic, and no variety can account for the incommensurability between lives and between same-size populations without extra theoretical resources.

The IER view comes equipped with the required theoretical resources. It diagnoses as the source of incommensurability the fact that some trade-offs are neither worth making nor worth not making and a small improvement fails to break the deadlock. 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 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 gumdrops and hangnails their proper axiological due. When a life is nearly blank, one fewer gumdrop and one extra hangnail can take it from good to bad. When a life is turbulent, gumdrops and hangnails pale almost into axiological insignificance. And because the IER view determines a life's categories of 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.⁴⁹

University of Oxford
elliott.thornley@philosophy.ox.ac.uk

REFERENCES

- Arrhenius, Gustaf. "Future Generations: A Challenge for Moral Theory." PhD diss., Uppsala University, 2000.
- . "An Impossibility Theorem for Welfarist Axiologies." *Economics and Philosophy* 16, no. 2 (October 2000): 247–66.
- Arrhenius, Gustaf, and Wlodek Rabinowicz. "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, 2015.
- . "Value Superiority." In *The Oxford Handbook of Value Theory*, edited by Iwao Hirose and Jonas Olson, 225–48. New York: Oxford University Press, 2015.
- Blackorby, Charles, Walter Bossert, and David Donaldson. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge: Cambridge University Press, 2005.
- . "Quasi-Orderings and Population Ethics." *Social Choice and Welfare* 13, no. 2 (1996): 129–50.
- Bossert, Walter. "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, forthcoming.
- Broome, John. "Loosening the Betterness Ordering of Lives: A Response to

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

- Rabinowicz." In *The Oxford Handbook of Population Ethics*, edited by Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns. Oxford: Oxford University Press, forthcoming.
- . *Weighing Goods*. Oxford: Blackwell, 1991.
- . *Weighing Lives*. Oxford: Oxford University Press, 2004.
- Bykvist, Krister. "The Good, the Bad, and the Ethically Neutral." *Economics and Philosophy* 23, no. 1 (March 2007): 97–105.
- Carlson, Erik. "Mere Addition and Two Trilemmas of Population Ethics." *Economics and Philosophy* 14, no. 2 (October 1998): 283–306.
- Chang, Ruth. "The Possibility of Parity." *Ethics* 112, no. 4 (July 2002): 659–88.
- De Sousa, Ronald B. "The Good and the True." *Mind* 83, no. 332 (October 1974): 534–51.
- Frick, Johann. "On the Survival of Humanity." *Canadian Journal of Philosophy* 47, nos. 2–3 (2017): 344–67.
- Gustafsson, Johan E. "Population Axiology and the Possibility of a Fourth Category of Absolute Value." *Economics and Philosophy* 36, no. 1 (March 2020): 81–110.
- Hudson, James L. "The Diminishing Marginal Value of Happy People." *Philosophical Studies* 51, no. 1 (January 1987): 123–37.
- Huemer, Michael. "In Defence of Repugnance." *Mind* 117, no. 468 (October 2008): 899–933.
- Meacham, Christopher J. G. "Person-Affecting Views and Saturating Counterpart Relations." *Philosophical Studies* 158, no. 2 (March 2012): 257–87.
- Nebel, Jacob M. "Totalism without Repugnance." In *Ethics and Existence: The Legacy of Derek Parfit*, edited by Jeff McMahan, Tim Campbell, James Godrich, and Ketan Ramakrishnan, 200–31. Oxford: Oxford University Press, 2021.
- Parfit, Derek. "Overpopulation and the Quality of Life." In *Applied Ethics*, edited by Peter Singer, 145–64. Oxford: Oxford University Press, 1986.
- . *Reasons and Persons*. Oxford: Clarendon Press, 1984.
- Qizilbash, Mozaffar. "The Mere Addition Paradox, Parity and Vagueness." *Philosophy and Phenomenological Research* 75, no. 1 (July 2007): 129–51.
- . "On Parity and the Intuition of Neutrality." *Economics & Philosophy* 34, no. 1 (2018): 87–108.
- Rabinowicz, Wlodek. "Broome and the Intuition of Neutrality." *Philosophical Issues* 19, no. 1 (October 2009): 389–411.
- . "Getting Personal: The Intuition of Neutrality Reinterpreted." In *Studies on Climate Ethics and Future Generations*, vol. 2, edited by Paul Bowman and Katharina Berndt Rasmussen, 59–90. Stockholm: Institute for Future

- Studies, 2020. <https://www.iffs.se/en/publications/working-papers/studies-on-climate-ethics-and-future-generations-vol-2/>.
- Raz, Joseph. *The Morality of Freedom*. Oxford: Oxford University Press, 1986.
- Tännsjö, Torbjörn. "Why We Ought to Accept the Repugnant Conclusion." *Utilitas* 14, no. 3 (November 2002): 339–59.
- Thornley, Elliott. "A Dilemma for Lexical and Archimedean Views in Population Axiology." *Economics and Philosophy*. Published ahead of print, September 6, 2021. <https://doi.org/10.1017/s0266267121000213>.