

*The Disjunctive Repugnant Conclusion**

Johan E. Gustafsson[†]

ABSTRACT. Critical-Level Utilitarianism entails one of the Repugnant Conclusion and the Sadistic Conclusion (both of which are counter-intuitive), depending on the critical level. Indeterminate Critical-Level Utilitarianism is a version of Critical-Level Utilitarianism where it is indeterminate which well-being level is the critical level. Undistinguished Critical-Range Utilitarianism is variant of Critical-Level Utilitarianism where additions of lives in a range of well-being between the good and the bad lives makes the resulting outcome incomparable to the original outcome. These views both avoid the Repugnant Conclusion and avoid the Sadistic Conclusion. And they agree about all comparisons of outcomes that do not involve indeterminacy or incomparability. So it is unclear whether we have any reason to favour one of these theories over the other. I argue that Indeterminate Critical-Level Utilitarianism still entails the disjunction of the Repugnant Conclusion and the Sadistic Conclusion, which is also repugnant. Whereas, Undistinguished Critical-Range Utilitarianism does not entail this conclusion.

A central challenge in population axiology is to find a plausible way to avoid the following, counter-intuitive, conclusions:

The Repugnant Conclusion Each outcome where everyone has very good well-being is worse than some outcome where everyone has barely good well-being.¹

The Sadistic Conclusion Each outcome where everyone has bad well-being is better than some outcome where everyone has good well-being.²

While it's easy to find seemingly compelling axiologies that avoids one of these conclusions, it's surprisingly hard to find a compelling axiology

* Forthcoming in *Utilitas*.

[†] I would be grateful for any thoughts or comments on this paper. They can be sent to me at johan.eric.gustafsson@gmail.com.

¹ Parfit's (1984, p. 388) canonical statement of the Repugnant Conclusion is more complicated. But his complications won't matter for the present discussion.

² Arrhenius 2000, p. 256.

that avoids both.³ Consider, for instance, *Total Utilitarianism*, the axiology which values an outcome by the sum total of the well-being of the individuals who exist in the outcome. While Total Utilitarianism avoids the Sadistic Conclusion, it entails the Repugnant Conclusion.⁴

Critical-Level Utilitarianism is a generalization of Total Utilitarianism. Critical-Level Utilitarianism, just like Total Utilitarianism, is an additive measure of the value of an outcome, summing up the well-being of each person to get the value of the outcome. The difference is that, according to Critical-Level Utilitarianism, a certain critical level is subtracted from each individual's well-being level before these are summed up. If this critical level is the neutral level of well-being, however, the resulting version of Critical-Level Utilitarianism is equivalent to Total Utilitarianism. Let the critical-total value of an outcome X given a certain critical level w be

$$v(X, w) = \sum_{i \text{ exists in } X} (w_{i,X} - w),$$

where $w_{i,X}$ is the well-being of individual i in X . Then, according to


Critical-Level Utilitarianism An outcome X is at least as good as an outcome Y if and only if $v(X, w)$ is at least as great as $v(Y, w)$.⁵

No matter which level is the critical level, either Critical-Level Utilitarianism entails the Repugnant Conclusion or it entails the Sadistic Conclusion. To see this, note first that, given a neutral or bad critical level, Critical-Level Utilitarianism entails the Repugnant Conclusion. Let A be any outcome where everyone has very good well-being. Then an outcome Z with lives that are barely good that will have a greater critical total value given that it is sufficiently large:

³ While we will explore how the Repugnant Conclusion can be avoided, it's not clear that it needs to be avoided. The intuition that the Repugnant Conclusion is false may be misleading. See Gustafsson 2022 and Zuber et al. 2021.


⁴ For example, while Average Utilitarianism avoids the Repugnant Conclusion, it entails the Sadistic Conclusion. See Arrhenius 2000, p. 251. See Parfit 1984, p. 422 for the similar objection that, on Average Utilitarianism, it can be better to add lives at a negative well-being level than to not do so (other things being equal).

⁵ Blackorby and Donaldson 1984, p. 14. Broome (2004, p. 255) calls this the Integrated Standardized Total Principle.



To avoid the Repugnant Conclusion, we need a critical level sufficiently above the neutral level so that it is unrepugnant that an outcome with a very large population where everyone has a level of well-being barely above the critical level is better than an outcome with a smaller population where everyone has very good well-being.

Yet, if the critical level is above the neutral level with some margin so that there is at least one good well-being level below the critical level, Critical-Level Utilitarianism entails the Sadistic Conclusion, which seems at least as repugnant as the Repugnant Conclusion. The Sadistic Conclusion follows from Critical-Level Utilitarianism if there is a good well-being level below the critical level. To see this, note that one could make an arbitrarily bad outcome by increasing the size of a population where everyone has this good level of well-being below the critical level. Let Bad A be any outcome where everyone has bad well-being. Then, like before, there will be an outcome Z consisting of lives that are barely good that will have a lower critical total value given that it is sufficiently large:



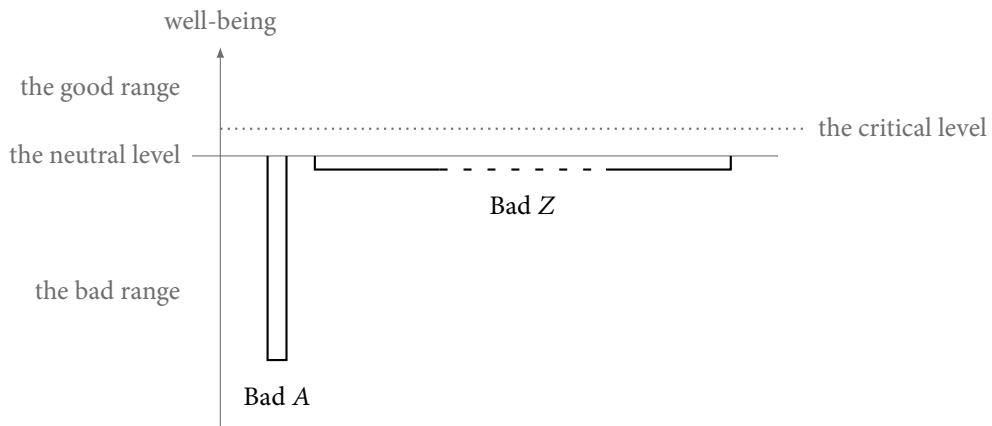
So, to avoid the Sadistic Conclusion, we need a bad or almost neutral

critical level.

There is also a negative analogue to the Repugnant Conclusion. In this variant, people's well-being is bad rather than good. According to

The Mirrored Repugnant Conclusion Each outcome where everyone has very low bad well-being is better than some outcome where everyone has barely bad well-being.⁶

Unless the critical level is sufficiently lower than the neutral level, Critical-Level Utilitarianism entails the Mirrored Repugnant Conclusion.



The trouble is that the critical level needs to be sufficiently above the neutral level to avoid the Repugnant Conclusion. So, once we specify the critical level, Critical-Level Utilitarianism cannot avoid either entailing the Repugnant Conclusion or entailing both of the Mirrored Repugnant Conclusion and the Sadistic Conclusion.

1. Indeterminate Critical-Level Utilitarianism

One way to amend Critical-Level Utilitarianism so that it avoids each of these counter-intuitive conclusions is to add that it is indeterminate which well-being level is the critical level.

⁶ Carlson 1998, p. 297.

Indeterminate Critical-Level Utilitarianism is Critical-Level Utilitarianism in conjunction with there being a range of indeterminacy, which is a range of well-being levels such that it is determinate that the critical level is in this range but, for each level in the range, it is indeterminate whether that level is the critical level.

This makes some assumptions about the logic of indeterminacy. We have a determinately true disjunction where no disjunct is determinately true, that is, the disjunction of the claims, for each well-being level in the critical range, that that level is critical level.⁷

John Broome's account of Indeterminate Critical-Level Utilitarianism relies on supervaluationism.⁸ According to supervaluationism, the evaluation of a statement involving vague or indeterminate terms depends on its sharpenings, which are the admissible ways it can be made precise. On Broome's version of supervaluationism, we may assert a statement 'S' if and only if 'S' is true under each one of its sharpenings.⁹ To apply supervaluationism to Indeterminate Critical-Level Utilitarianism, we let the well-being levels in the range of indeterminacy be the extensions of the sharpenings of the term 'critical level', and we assume for simplicity that the indeterminacy of 'critical level' is the only source of indeterminacy. We can then assert the statement 'one of the well-being levels in the range of indeterminacy is the critical level', since this will be true given every sharpening of 'critical level'. But, for every well-being level *l* in the range of indeterminacy, we may not assert the statement '*l* is the critical level', since this will be false given some sharpening of 'critical level' (assuming that there is more than one level in the range).

Indeterminate Critical-Level Utilitarianism could also be coupled with an epistemic view of indeterminacy, where indeterminacy is merely a kind of ignorance. On the epistemic view, there are sharp boundaries to the extensions of vague or indeterminate terms but we do not know

⁷ While determinately true disjunctions where no disjunct is determinately true are allowed on supervaluationism and an epistemic view like Williamson's (1994), they are ruled out by truth-functional many-valued theories such as Halldén's (1949, p. 48) logic of nonsense and Machina's (1976) degree-theoretic approach, where a disjunction cannot be completely true if no disjunct is completely true.

⁸ Broome 2004, pp. 255–6.

⁹ Broome (2004, pp. 177–8) adds the qualification that supervaluationism does not apply if 'S' contains terms from the metalanguage such as, for example, 'vagueness', 'sharpening', or 'we may assert'.

and perhaps cannot know where those boundaries lie.¹⁰ Applying the epistemic view to Indeterminate Critical-Level Utilitarianism, we let the well-being levels in the range of indeterminacy be a set of well-being levels such that we know that one of them is the extension of ‘critical level’ and, if there is more than one level in the range, we are ignorant for each level in the range whether that level is the extension of ‘critical level’. So we know that one of the well-being levels in the range of indeterminacy is the critical level. But, for every well-being level in the range of indeterminacy, we are ignorant about whether that level is the critical level (assuming again that there is more than one level in the range).

Let I be the set of well-being levels in the range of indeterminacy. Then, on either supervalueationism or the epistemic view of indeterminacy, Indeterminate Critical-Level Utilitarianism entails that

an outcome X is determinately at least as good as an outcome Y if and only if, for all w in I , it holds that $v(X, w)$ is at least as great as $v(Y, w)$, and

an outcome X is determinately not at least as good as an outcome Y if and only if, for all w in I , it holds that $v(X, w)$ is not at least as great as $v(Y, w)$.

Indeterminate Critical-Level Utilitarianism avoids each one of the Repugnant Conclusion, the Mirrored Repugnant Conclusion, and the Sadistic Conclusion. Given a range of indeterminacy with an upper bound sufficiently higher than the neutral level and a lower bound sufficiently lower than the neutral level, none of these conclusions is determinately true on Indeterminate Critical-Level Utilitarianism. Yet, while these conclusions are not determinately true on Indeterminate Critical-Level Utilitarianism, they are not determinately false either. We might feel, however, following Broome, that the fact that these conclusions aren’t determinately true is enough to curb their intuitive repugnance.¹¹

2. Undistinguished Critical-Range Utilitarianism

We have seen how we can avoid both the Repugnant Conclusion and the Sadistic Conclusion if we introduce indeterminacy to Critical-Level Utilitarianism.

¹⁰ Keefe and Smith 1997, p. 17.

¹¹ Broome 2004, pp. 213–14.

tarianism. But, in much the same way, we can avoid these conclusions if we instead introduce incomparability to the theory. There is a structurally analogous variation of Critical-Level Utilitarianism, where, instead of a range of indeterminacy there is range of undistinguishedness — a range of well-being levels such that adding a person with a life at those levels of well-being makes the resulting outcome incomparable to the original outcome. Consider

Critical-Range Utilitarianism Let U be the set of well-being levels in the range of undistinguishedness. Then an outcome X is at least as good as an outcome Y if and only if, for all w in U , it holds that $v(X, w)$ is at least as great as $v(Y, w)$.¹²

If U and I contain the same well-being levels, then Critical-Range Utilitarianism will entail that an outcome X is at least as good as an outcome Y if and only if Indeterminate Critical-Level Utilitarianism entails that X is determinately as least as good as Y . The theories still differ, however, in that, when Indeterminate Critical-Level Utilitarianism entails that X is indeterminately as least as good as Y , Critical-Range Utilitarianism entails that X is incomparable with Y .

Suppose we further accept that lives at the well-being levels in U are *undistinguished* — that is, that they are value bearers but they are not good, not bad, and not neutral.¹³ Then we can accept

The Personal Critical-Range View Let U be the set of well-being levels in the range of undistinguishedness. Then lives at well-being levels higher than all levels in U are good, lives at well-being levels lower than all levels in U are bad, and lives at well-being levels in U are undistinguished.¹⁴

Note that, since undistinguished lives are still value bearers, this view is consistent with the claim that a first life is at least as good as a second life if and only if the first life is at a well-being level at least as high as the second life. That is, we can compare lives at well-being levels in the range of undistinguishedness.

¹² Blackorby et al. (1996, p. 141) called this Incomplete Critical-Level Utilitarianism. They (2005, p. 252) later renamed it Critical-Band Utilitarianism. Rabinowicz (2009, p. 404), Qizilbash (2007), and Gustafsson (2020, p. 92) propose a similar views.

¹³ Gustafsson 2020, p. 88; 2023, p. 610.

¹⁴ See Rabinowicz 2009, p. 391; 2022, pp. 129–30 and Gustafsson 2020, p. 87.

Let *Undistinguished Critical-Range Utilitarianism* be the combination of Critical-Range Utilitarianism and the Personal Critical-Range View.¹⁵ Given Undistinguished Critical-Range Utilitarianism, we avoid both the Repugnant Conclusion and the Sadistic Conclusion (and the Mirrored Repugnant Conclusion). So Indeterminate Critical-Level Utilitarianism and the combination of Critical-Range Utilitarianism and the Personal Critical-Range View both avoid repugnance and sadism.¹⁶

There are, of course, other worries. For instance, Critical-Range Utilitarianism is open to the *Greediness Objection* — that is, if we improve the lives in an outcome and add some people with a well-being level in the critical range, the resulting outcome may be incomparable to the original. One motivation for having a critical range is to capture the *Intuition of Neutrality* — that is, to capture the intuition that, while we are in favour of making people happy, we are neutral about creating happy people.¹⁷ But, if improving people's lives is an improvement and adding a life is neutral, then the combination should result in an outcome that is better, rather than an outcome that is incomparable with the original.¹⁸ Nevertheless, Indeterminate Critical-Level Utilitarianism faces much the same problem, since, if we improve the lives in an outcome and add some people with a well-being level in the critical range, the resulting outcome whose comparison to the original may be indeterminate according to Indetermi-

¹⁵ Gustafsson 2020, p. 92.

¹⁶ It may be objected that Undistinguished Critical-Range Utilitarianism entails both of the following conclusions (and their good/bad-inverted, mirrored variants):

The Weak Repugnant Conclusion (barely-not-bad version) Each outcome where everyone has very good well-being is not at least as good as some outcome where everyone is at a well-being level that is just marginally higher than some bad well-being level. (Gustafsson 2020, pp. 95.)

The Weak Sadistic Conclusion (barely-not-bad version) Each outcome where everyone has bad well-being levels is not worse than some outcome where everyone is at a well-being level that is just marginally higher than some bad well-being level.

These conclusions, however, seem less repugnant once we grant the Personal Critical-Range View. An outcome with a very large number of lives at an undistinguished level will be (overall) very undistinguished and, hence, incomparable with many outcomes with only bad lives or only good lives. For a fuller response along this line, see Gustafsson 2020, pp. 94–6.

¹⁷ Narveson 1973, p. 80 and Broome 2004, pp. 143–6.

¹⁸ Broome 2004, p. 169. For a repose to the Greediness Objection, see Gustafsson 2020, pp. 107–10.

nate Critical-Level Utilitarianism — rather than determinately better.¹⁹

3. The Disjunctive Repugnant Conclusion

So we have two approaches for amending Critical-Level Utilitarianism with essentially the same formal structure. And they both avoid entailing the various problematic conclusions we have considered. But they differ in whether this is done by introducing indeterminacy or incomparability. As Broome notes, indeterminacy and incomparability lead to much the same practical problems.²⁰ And the views agree about all comparisons of outcomes that do not involve indeterminacy or incomparability. That is, they will do so if the range of indeterminacy, I , includes the same well-being levels as the range of undistinguishedness, U . The two theories agree, for any two outcomes X and Y , about whether X is determinately at least as good as Y .²¹ (And, the outcomes are incomparable according to Critical-Range Utilitarianism if and only if their comparison is indeterminate on Indeterminate Critical-Level Utilitarianism.) So one may wonder if we have any reason to favour one of these approaches over the other.

I shall argue, however, that these approaches differ in their ability to avoid repugnance. Even if we merely wish to avoid entailing that something repugnant is *determinately* true, Indeterminate Critical-Level Utilitarianism still fails in regard. While the two theories agree, for any two outcomes X and Y , whether X is determinately at least as good as Y , the theories can still disagree about general claims about comparisons of outcomes. To see this, we shall consider two conclusions which are disjunctions of the previous ones. According to

¹⁹ Rabinowicz 2009, p. 400.

²⁰ Broome 2004, p. 85; 2009, pp. 415–17.

²¹ For example, suppose the range of both indeterminacy and undistinguishedness is $[-1, 1]$. And suppose that we are considering adding a pair of people at well-being levels $(1000, 0.99)$. Then, on Indeterminate Critical-Level Utilitarianism, this is a determinately good addition — since, for every sharpened critical level w in the range of indeterminacy, it holds that $(1000 - w) + (0.99 - w) > 0$. But, on Undistinguished Critical-Range Utilitarianism, this is also a good addition — since, for every level w in the range of undistinguishedness, it holds that $(1000 - w) + (0.99 - w) > 0$. Note that the addition is overall good, even though the addition of the person at level 0.99 is undistinguished, since, on every level w in the range of undistinguishedness, the contribution of the person at level 1000 outweighs the (potentially) negative contribution of the person at level 0.99.

The Disjunctive Repugnant Conclusion [Either the Repugnant Conclusion or the Mirrored Repugnant Conclusion] is true.²²

And, similarly, according to

The Either-Repugnant-or-Sadistic Conclusion [Either the Repugnant Conclusion or the Sadistic Conclusion] is true.

These two disjunctive conclusions strike me as on a par in repugnance with the Repugnant Conclusion.

Moreover, there is a straightforward argument that the Disjunctive Repugnant Conclusion should be repugnant. For each disjunct in these conclusions, any value ordering of outcomes that satisfies the disjunct is repugnant. Hence each of these disjuncts entails that the value ordering of outcomes is repugnant. By argument by cases, we then have that each of these disjunctions entails that the value ordering of outcomes is repugnant. So these disjunctive conclusions should also be repugnant. Formally, the argument can be put as follows. Let RC be the Repugnant Conclusion, let SC be the Sadistic Conclusion, and let R be that the value ordering of outcomes is repugnant. Since any value ordering of outcomes satisfying at least one of RC and SC is repugnant, we have $RC \supset R$ and $SC \supset R$. Then from $RC \supset R$, $SC \supset R$, and $RC \vee SC$, we have R by argument by cases.²³

There is a further argument that these disjunctive conclusions should be repugnant if ‘repugnant’ is used in an evaluative sense, which is how it is used by, for example, J. M. E. McTaggart, who was the first to call conclusions like the Repugnant Conclusion ‘repugnant’. He writes:

²² Here and in the Either-Repugnant-or-Sadistic Conclusion, the square brackets indicate that ‘is true’ applies to the disjunction rather than the disjuncts.

²³ It may be objected that argument by cases is not a valid form of inference on supervaluationism, see Williamson 1994, p. 152. Keefe (2000, pp. 175–6), however, points out that, in the absence of the ‘definitely’ operator, the consequence relation of supervaluationism coincides with the consequence relation of classical logic. And there is no ‘definitely’, ‘determinately’, or ‘we may assert’ operator in play in the present argument — there is, in fact, no talk of vagueness or indeterminacy at all. In the formal version of the argument, we also avoid the ‘is true’ locution of the disjunctive conclusions. Hence, in this case, we can apply classical inference rules such as argument by cases even on supervaluationism. Furthermore, argument by cases is valid on Williamson’s (1994, p. 271) logic of clarity.

This conclusion would, I believe, be repugnant to certain moralists. But, in the first place, a conclusion may be rightly repugnant to us, and yet it may be true, since the universe is not completely good.²⁴

A standard view in the logic of value is the principle of disjunctive interpolation, that is, the principle that disjunctions fall between their disjuncts in terms of intrinsic value.²⁵ If, for example, p is intrinsically bad and q is intrinsically bad, then the disjunction $p \vee q$ is also intrinsically bad. Analogously, the repugnance of a disjunction should fall between its disjuncts in terms of repugnance. Hence, if both the Repugnant Conclusion and the Mirrored Repugnant Conclusion are repugnant, the Disjunctive Repugnant Conclusion should also be repugnant. Likewise, if both the Repugnant Conclusion and the Sadistic Conclusion are repugnant, the Either-Repugnant-or-Sadistic Conclusion should also be repugnant.

The problem is that, regardless of which well-being levels are in the range of indeterminacy, Indeterminate Critical-Level Utilitarianism entails both of these disjunctive conclusions. Or, at least, it does so given either of the two standard views of indeterminacy, supervenience and

²⁴ McTaggart 1927, p. 453.

²⁵ See, for example, Hansson 1968, p. 438, Castañeda 1969, p. 264, and Chisholm 1975, p. 297. The name is due to Hansson (1997, p. 437). This principle has been challenged by Oldfield (1977, pp. 245–6). For Oldfield's counter-example, we assume hedonism and consider all conjunctions such that the first conjunct is the state of affairs Jones's being happy to degree 1 and the second conjunct is either the state that Bill does not exist or a state of the form Bill's being happy to degree i , where positive values of i represents degrees of happiness and negative values of i represents degrees of unhappiness. Some of these conjunctions are intrinsically good; some are intrinsically bad; and some are intrinsically neutral. Let G be the disjunction of the intrinsically good conjunctions; let B be the disjunction of the intrinsically bad conjunctions; and let N be the disjunction of the intrinsically neutral conjunctions. It seems plausible that G is intrinsically good, B is intrinsically bad, and N is intrinsically neutral. Oldfield claims that the disjunction G -or- B is neither intrinsically good nor intrinsically bad, since it is 'good at some worlds and bad at others'. Finally, on hedonism, the disjunction N -or- $(G$ -or- $B)$ must be intrinsically good, since it is necessarily equivalent with Jones's being happy to degree 1. But then N -or- $(G$ -or- $B)$ is intrinsically better than each of its disjuncts, which contradicts the principle of disjunctive interpolation. A weak point in Oldfield's counter-example is his argument that G -or- B is not intrinsically good. It turns on Oldfield's idea that intrinsic value can be world relative, which, as Carlson (1997, p. 99) points out, seems to 'reflect a failure to grasp the concept of intrinsic value'. Carlson (1997, p. 104), however, still finds it plausible that G -or- B isn't intrinsically good and he endorses Oldfield's counter-example. Yet I don't see why G -or- B cannot be intrinsically good; it guarantees that Jones is happy to degree 1 and it does not guarantee that there is any unhappiness.

the epistemic view.²⁶ Let a *precise version of Critical-Level Utilitarianism* be Critical-Level Utilitarianism with a certain well-being level as the precise critical level. So, for each well-being level, we have a unique precise version of Critical-Level Utilitarianism where that level is the critical level. Every precise version of Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion. This is because, whichever well-being level is the critical level, Critical-Level Utilitarianism entails at least one of the disjuncts in the Disjunctive Repugnant Conclusion. Moving on, the rest of the argument proceeds differently depending on whether we adopt supervaluationism or the epistemic view of indeterminacy.

On supervaluationism, we may assert the statement ‘Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion’ if it is true on each one of its sharpenings. Having assumed for simplicity that the only source of indeterminacy is the indeterminacy of ‘critical level’, each sharpening of the above statement will state that a certain precise version of Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion. And, since each precise version of Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion, we have that each sharpening of the above statement is true. Hence we may assert that Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion, even if the critical level is indeterminate.

On the epistemic view, we know that one well-being level in the range of indeterminacy is the critical level. So, if we know that Indeterminate Critical-Level Utilitarianism is true, we know that some precise version of Critical-Level Utilitarianism is true, even though we do not know which version is true. But, since every precise version of Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion, we would then know that the Disjunctive Repugnant Conclusion is true. Hence we know that Indeterminate Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion.

So, on both of these views, we find that Indeterminate Critical-Level Utilitarianism entails the Disjunctive Repugnant Conclusion. And, changing what needs to be changed, we find that Indeterminate Critical-Level Utilitarianism entails the Either-Repugnant-or-Sadistic Conclusion. Thus Indeterminate Critical-Level Utilitarianism does not

²⁶ Also, it is hard to see how to make Indeterminate Critical-Level Utilitarianism fit with some of the less standard views of indeterminacy. For example, as was mentioned in note 7, truth-functional many-valued theories do not fit with the assumptions in Indeterminate Critical-Level Utilitarianism about the logic of indeterminacy.

avoid these repugnant disjunctive conclusions.

The Undistinguished Critical-Range Utilitarianism, however, entails neither Disjunctive Repugnant Conclusion nor the Either-Repugnant-or-Sadistic Conclusion. It entails both of these conclusions are false, since it entails that each of their disjuncts are false. Hence we have an argument for favouring Undistinguished Critical-Range Utilitarianism over the formally very similar Indeterminate Critical-Level Utilitarianism.

I wish to thank Gustaf Arrhenius, Barry Lee, Martin Peterson, Mozaffar Qizilbash, Wlodek Rabinowicz, and Tom Stoneham for valuable comments.

References

- Arrhenius, Gustaf (2000) 'An Impossibility Theorem for Welfarist Axiologies', *Economics and Philosophy* 16 (2): 247–266.
- Blackorby, Charles, Walter Bossert, and David Donaldson (1996) 'Quasi-Orderings and Population Ethics', *Social Choice and Welfare* 13 (2): 129–150.
- (2005) *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*, Cambridge: Cambridge University Press.
- Blackorby, Charles and David Donaldson (1984) 'Social Criteria for Evaluating Population Change', *Journal of Public Economics* 25 (1): 13–33.
- Broome, John (2004) *Weighing Lives*, Oxford: Oxford University Press.
- (2009) 'Reply to Rabinowicz', *Philosophical Issues* 19 (1): 412–417.
- Carlson, Erik (1997) 'The Intrinsic Value of Non-Basic States of Affairs', *Philosophical Studies* 85 (1): 95–107.
- (1998) 'Mere Addition and Two Trilemmas of Population Ethics', *Economics and Philosophy* 14 (2): 283–306.
- Castañeda, Hector-Neri (1969) 'Ought, Value, and Utilitarianism', *American Philosophical Quarterly* 6 (4): 257–275.
- Chisholm, Roderick M. (1975) 'The Intrinsic Value in Disjunctive States of Affairs', *Noûs* 9 (3): 295–308.
- Gustafsson, Johan E. (2020) 'Population Axiology and the Possibility of a Fourth Category of Absolute Value', *Economics and Philosophy* 36 (1): 81–110.
- (2022) 'Our Intuitive Grasp of the Repugnant Conclusion', in Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns, eds., *The Oxford Handbook of Population Ethics*, chap. 16, pp. 371–389, New York: Oxford University Press.

- (2023) ‘The Possibility of Undistinguishedness,’ *Philosophical Studies* 180 (2): 609–613.
- Halldén, Sören (1949) *The Logic of Nonsense*, Uppsala Universitets Årsskrift: Uppsala.
- Hansson, Bengt (1968) ‘Fundamental Axioms for Preference Relations,’ *Synthese* 18 (4): 423–448.
- Hansson, Sven Ove (1997) ‘Situationist Deontic Logic,’ *Journal of Philosophical Logic* 26 (4): 423–448.
- Keefe, Rosanna (2000) *Theories of Vagueness*, Cambridge: Cambridge University Press.
- Keefe, Rosanna and Peter Smith (1997) ‘Introduction: Theories of Vagueness,’ in Rosanna Keefe and Peter Smith, eds., *Vagueness: A Reader*, pp. 1–57, Cambridge, MA: MIT Press.
- Machina, Kenton F. (1976) ‘Truth, Belief, and Vagueness,’ *Journal of Philosophical Logic* 5 (1): 47–78.
- McTaggart, John McTaggart Ellis (1927) *The Nature of Existence Volume II*, Cambridge: Cambridge University Press.
- Narveson, Jan (1973) ‘Moral Problems of Population,’ *Monist* 57 (1): 62–86.
- Oldfield, Edward (1977) ‘An Approach to a Theory of Intrinsic Value,’ *Philosophical Studies* 32 (3): 233–249.
- Parfit, Derek (1984) *Reasons and Persons*, Oxford: Clarendon Press.
- Qizilbash, Mozaffar (2007) ‘The Mere Addition Paradox, Parity and Vagueness,’ *Philosophy and Phenomenological Research* 75 (1): 129–151.
- Rabinowicz, Wlodek (2009) ‘Broome and the Intuition of Neutrality,’ *Philosophical Issues* 19 (1): 389–411.
- (2022) ‘Getting Personal: The Intuition of Neutrality Reinterpreted,’ in Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns, eds., *The Oxford Handbook of Population Ethics*, pp. 114–141, New York: Oxford University Press.
- Williamson, Timothy (1994) *Vagueness*, London: Routledge.
- Zuber, Stéphane, Nikhil Venkatesh, Torbjörn Tännsjö, Christian Tarsney, H. Orri Stefánsson, Katie Steele, Dean Spears, Jeff Sebo, Marcus Pivato, Toby Ord, Yew-Kwang Ng, Michal Masny, William MacAskill, Nicholas Lawson, Kevin Kuruc, Michelle Hutchinson, Johan E. Gustafsson, Hilary Greaves, Lisa Forsberg, Marc Fleurbaey, Diane Coffey, Susumu Cato, Clinton Castro, Tim Campbell, Mark Budolfson, John Broome, Alexander Berger, Nick Beckstead, and Geir B. Asheim (2021) ‘What Should We Agree on about the Repugnant Conclusion?’, *Utilitas* 33 (4): 379–383.