Specifying Welfare Functions according to Unreal and Trivial Boundary Conditions

Walter Barta

University of Houston

DRAFT Fall, 2018

(as presented at the 2018 HCC Humanities and Philosophy Conference)

0 Introduction

In his book *Reasons and Persons*, Derek Parfit proposes the search for a self-consistent theory of population ethics, a theory capable of answering questions about the welfares of populations in a manner that satisfies all of our ethical intuitions, what he calls "Theory X.".[1] But in the same work, Parfit offers what he sees as a major obstacle to that goal, the so-called "Repugnant Conclusion", worrying whether the most well-off population is an increasingly large population. This problem, along with Roderick Ninian Smart's "Negative Utilitarianism" and Robert Nozick's "Utility Monster" belong to a special class of ethical cases dealing with the mathematical limits of zero and infinity: trivial and unreal ethical solutions.[2][3][4] These kinds of problems have plagued population ethics since Thomas Malthus first recognized the problem of exponential growth and its deleterious effects on personal wellbeing.[5] In an attempt to answer Parfit, some, such as T. Sider, T. Hurka, and Y. Ng, have suggested formulations of marginal wellbeing without undermining core ethical principles. Other authors, like G. Arrhenius et al., have suggested that Parfit's problem fundamentally undermines the possibility of a self-consistent population ethics.[6]

Here we will suggest that these problems result from an improper inclusion of trivial and unreal numbers into the ethical domain, making what Parfit calls "mistakes in moral mathematics," and are thus resolved by an exclusion of trivial and unreal numbers from the ethical domain by carefully specifying a welfare function to exclude them.[7] Our analysis will consist of the following steps:

1. First, we will look at the ethical domain as containing a function of personal wellbeings dependent upon the variables of material goods and populations of persons.
2. Second, we will use ethically absurd cases as the boundary conditions within which real ethical solutions may be mathematically defined.

---

[1] Parfit, Derek, *Reasons and Persons*, Oxford: Clarendon Press, (Oxford, 1984). p. 390.
[2] Smart, R. N., "Negative Utilitarianism", *Mind*, 67: (1958), pp. 542–3.
[3] Nozick, *Robert. Anarchy, State, and Utopia*. Blackwell Publishers Ltd., (1999).
[4] Kant, Immanuel. *Critique of Pure Reason*. A548/B576. p. 473.
[5] Malthus, Thomas, Essay on the Principle of Population, (London, Penguin Books: 2015).
[6] Arrhenius, Gustaf, Ryberg, Jesper and Tännsjö, Torbjörn, "The Repugnant Conclusion", *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/repugnant-conclusion/>.
[7] Parfit, p 67.

3. Third, using such methods, we might propose some design specifications in the form of zero-value conditions that narrow down the ethical domain to exclude absurd cases.
4. Fourth, within those specifications, from within the subset of the real ethical domain, we can specify continuous, well-defined polynomial approximation functions that demonstrate the real possibility of such welfare functions.

We will conclude with a set of polynomial approximation candidates for a Theory X and defend its implications.

1 Ethical Domains

To characterize a self-consistent theory of population ethics, we must first establish the existence of an ethical domain, specifically a domain of possible states of affairs with considerations of the total wellbeing of possible populations of persons amidst varying supplies of material goods. (Note: the following analysis will use the notation: U Total wellbeing, u_i Individual wellbeing, u_m Marginal wellbeing, N Population, G Individual Goods, G_tTotal Goods.)

> **Ethical Domain**: For all states of affairs x there is some total wellbeing U which is a function of the population of persons N greater than or equal to zero and material goods G than or equal to zero, which is equal to the total wellbeing of that state of affairs.

As a corollary to the ethical domain condition, we can suppose that some suffering exists. Thus, we might stipulate that there are some states of affairs that are ethical problems, offering neutral or negative wellbeing.

> **Ethical Problems**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N greater than or equal to zero and material goods G than or equal to zero, which is less than zero.

As another corollary to the ethical domain condition, we might stipulate that there are some good scenarios. These states of affairs that would satisfy as ethical solutions, some states of affairs in the ethical domain that offer positive wellbeing.

> **Ethical Solutions**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N greater than or equal to zero and material goods G than or equal to zero, which is greater than zero.

We might further stipulate a subset of ethical solutions that are best case scenarios. These states of affairs that would satisfy as ethical optimums, some states of affairs in the ethical domain that offer the best wellbeing.

> **Ethical Optimums**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N greater than or equal to zero and material goods G than or equal to zero, which is maximum wellbeing.

With these bases established, we might consider their potential pitfalls.

2 Trivial and Unreal Ethical Cases

Two major problems emerge from this formulation of Ethical Domains: trivial ethical solutions and unreal ethical solutions, solutions involving zeros and infinities in the domain (Appendix A).

Population ethics notably fails to exclude these possibilities. The classical formulation of this is the utilitarian principle of Jeremy Bentham's "the greatest good for the greatest number", expresses no formal exclusion of unreal solutions, because it proposes no upper bounds to either "good" or "number", but only a range of "greatest" solutions on a plane of population and individual wellbeing (Appendix B: Table 2). "Greatest" is vague enough that the range of possible utilitarian outcomes could include solutions that are infinite, trivial, and non-unique.[8] In this sense, utilitarianism is an underspecified ethical system. In as much, the classical utilitarian function has many ethically and mathematically unappealing properties.

Derek Parfit suggests that only a critical examination of the fundamental assumptions of classical utilitarianism can ever resolve these trivial and unreal solutions because the problem is intrinsic to classical utilitarianism itself:

> We should not try to avoid this conclusion by appealing to principles covering some different part of morality. This conclusion is *intrinsically* repugnant. And this conclusion is implied by the Impersonal Total Principle, which is a particular version of the Principle of Beneficence….[9]

In a strong sense, the mathematical problems lie internal to Bentham's utilitarian formulation and must be solved with an appeal to principles fundamental to utilitarianism itself.

Some solutions to such problems have been offered, but many cause more dilemmas than they solve.[10] Promising avenues have arisen within the horizon of utilitarian thought. Some authors have modified the basic Bentham principle with theories of critical level.[11][12] But, critical-level theories fail to anticipate the ethical gap opened by the "Sadistic Conclusion": some populations exist in which it is better to add persons of negative, not positive wellbeing.[13] Another version of similar modification is found by delving deeper into the fundamental principle of "Marginal Utility" (here we treat "utility" and "wellbeing" interchangeably). Various authors have noted that the theory of marginal utility assumed in the premise will change the ethical theory that results in the conclusion, and is thus a fruitful point of departure, as has been suggested.[14][15][16] Variants on marginal utility have been offered: compromise act utilitarian theory, the variable-value view, number-dampened total utility. These versions offer promising avenues towards a

---

[8] Bentham, Jeremy. *An Introduction to the Principles of Moral and Legislation*. Anodos Books, (2019), p. 7.
[9] Parfit, p. 390.
[10] Arrhenius, Gustaf, Ryberg, Jesper and Tännsjö, Torbjörn, "The Repugnant Conclusion", *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/repugnant-conclusion/>.
[11] Broome, J., Weighing Lives, Oxford: Oxford University Press, (2004).
[12] Blackorby, C., W. Bossert and D. Donaldson, 2004, "Critical-Level Population Principles and the Repugnant Conclusion", *The Repugnant Conclusion: Essays on Population Ethics*. Edited by J. Ryberg and T. Tännsjö (eds.) (2004), pp. 45-60.
[13] Arrhenius, G.. "An Impossibility Theorem for Welfarist Axiology", *Economics and Philosophy*, 16: 247–266.
[14] Hurka, T., "Value and Population Size", *Ethics*, 93: (1983), (2000), pp. 496–507.
[15] Ng, Y.-K.. "What Should We Do About Future Generations? Impossibility of Parfit's Theory X", *Economics and Philosophy*, 5: (1989), pp. 135–253.
[16] Sider, T. R. "Might Theory X be a Theory of Diminishing Marginal Value?", *Analysis*, 51: (1991), pp. 265–71.

Theory X that do not abandon utilitarian principles altogether. However, none of these curves seems to satisfy all ethical intuitions. Perfecting a theory of Marginal Ueturns to fit our ethical intuitions might therefore be the most promising path forward. This might be seen as a movement towards aligning the "variable-value" view and the "critical-level" view with classical total utilitarianism by exploring various theories of Marginal Returns. However, all of these proposals offer solutions in tension with, not in conjunction with, Parfit's concept of Beneficence, positive wellbeing for persons, and thus are incomplete in the Ethical Domain.

A true solution would necessarily exclude these trivial and unreal solutions from the Ethical Domain. Real solutions to ethical problems would not be infinite, zero, or indeterminate; rather, real solutions would be finite, non-trivial, and unique. We can infer this from a slough of other well-regarded ethical principles. Per Immanuel Kant, so long as we accept the aphorism that "ought implies can", and as long as what we can do is what it is possible to do, we must specify impossible solutions out of ethical expectations.[17] In other words, any adequately specified ethical system would not admit of the unreal. Infinities, zeroes, and indeterminacies lie outside of the real domain of possibility, and yet still many ethical systems inadvertently or reluctantly admit of these unreal solutions. Inadvertently, this is also consistent with a type of golden mean virtue ethics: extremes are vices and means are virtues, much in the sense of Aristotelian ethics.[18] Aristotle identified this problem in a slightly different way, but a way that may be made compatible with Kant. To Aristotelian ethics, virtue is a mean and vice an extreme: "evil belongs to the class of the unlimited, as the Pythagoreans conjectured, and good to that of the limited" (Aristotle).[19] The unlimited, in other words unreal, solutions would need to be omitted from the range of reasonable ethical solutions. To Thomas Hobbes and his contractualist ethical position it was likewise axiomatic that infinities could not be rationally conceived, and therefore to Hobbes infinities could not be rationally expected of an ethical system.[20]

With these considerations in mind, we might attempt to consider trivial and unreal ethical solutions in turn and offer further bounds on the ethical domain that exclude them while still being consistent with the Principle of Beneficence. (Note that the following analysis is summarized formally in the table Appendix A: Proofs: Table 1, and the corresponding steps are provided parenthetically in the form S#.)

2.1 Trivial Cases

A trivial ethical optimum occurs when some ethical optimum has one or more of the functional variables set arbitrarily to zero.

> **Trivial Ethical Optimums**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N equal to zero and/or material goods G equal to zero, which is equal to the maximum wellbeing.

At least two trivial ethical optimums exist: 1) Lossless Depopulation, the case of zero population of persons; 2) Lossless Deprivation, the case of zero material goods.

---

[17] Parfit, p. 67.
[18] Bartlett, Robert C.; Collins, Susan D., *Nicomachean Ethics,* (Chicago: University of Chicago Press, 2011).
[19] Aristotle, ii.6.
[20] Hobbes, Thomas, *Leviathan*, (New York: Penguin Books, 1985), p. 99.

2.1.1 Lossless Depopulation

Description

The first type of trivial ethical optimum is Lossless Depopulation in the domain of population.

In his essay "Negative Utilitarianism", Roderick Ninian Smart suggests that perhaps the best world is one in which there are no suffering persons:

> Suppose that a ruler controls a weapon capable of instantly and painlessly destroying the human race. Now it is empirically certain that there would be some suffering before all those alive on any proposed destruction day were to die in the natural course of events. Consequently the use of the weapon is bound to diminish suffering, and would be the ruler's duty on [negative utilitarian] grounds.[21]

In other words, perhaps one way to reduce suffering is to eliminate suffering persons.

Definition

So, we may formally define Lossless Depopulation:

> **Lossless Depopulation**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N equal to zero and material goods G greater than or equal to zero, which is equal to the maximum wellbeing (S2.1.1.a).

We can also specify that this maximum must be indeterminate or zero because either nothing can be known or experienced about the wellbeing of a nonexistent being:

> **Indeterminate/Zero Nonbeing**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N equal to zero and material goods G greater than or equal to zero, which is indeterminate or zero (S2.1.1.b).

Furthermore, because nonbeing is both indeterminate or zero and the maximum wellbeing, it follows from the definition of a maximum that all other states of affairs must have well beings that are either indeterminate or zero. In other words, there is no existence better than non-existence. Even the most pampered and privileged person has a negative wellbeing. The entire ethical domain is negative.

> **Universally Indeterminate/Zero Wellbeing**: For all states of affairs x there is some total wellbeing U which is a function of the population of persons N greater than or equal to zero and material goods G than or equal to zero, which is either equal or less than indeterminate or zero (S2.1.1.c).

Rejection

However, we can reject Lossless Depopulation as trivial (S2.1.1), and thus specify it away, by making one assumption about the range of wellbeings, that there were even one determinate positive wellbeing for populations of persons greater than zero,

> **Incidentally Positive Wellbeing**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N greater than or equal to zero and material goods G than or equal to zero, which is determinate and greater than zero (S2.1.1.d).

---

[21] Smart, p. 542.

Barring certain kinds of pessimism and nihilism, it is intuitive to assume that Incidentally Positive Wellbeings exist: that at least one state of affairs has positive wellbeing (S2.1.1.d), and assuming this we have contradicted a Universally Indeterminate/Zero Wellbeing (S2.1.1.c), and thus Lossless Depopulation can be rejected.

2.1.2 Lossless Deprivation

Definition

The second type of trivial ethical optimum is Lossless Deprivation in the domain of material goods.

Ee may formally define Lossless Deprivation:

> **Lossless Deprivation**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N greater than or equal to zero and material goods G equal to zero, which is equal to the maximum wellbeing (S2.1.2.a).

Rejection

However, we can reject Lossless Deprivation as trivial (S2.1.2), and thus specify it away, by making one assumption about the materiality of persons, such that populations of persons require at least some goods as the substrate of their existence:

> **Material Person Condition**: For all states of affairs, if the material goods G equals to zero, then the population of persons N equals to zero (S2.1.2.b).

(Note: the Material Person Condition is not meant to exclude immaterial cases from the ethical domain, but it restricts the scope of the present analysis to the material domain, setting aside other cases for metaphysical consideration.)

Restricting oneself to considerations of material domains, if we assume Lossless Deprivation (S2.1.2.a) and the Material Person Condition holds (S2.1.2.b), then all cases of zero material goods must be cases of zero material populations of persons, so the Lossless Deprivation case implies Lossless Depopulation (S2.1.2.c). However, since we have already excluded Lossless Depopulation as false (S2.1.1), for being trivial, then Lossless Deprivation must be false, and can be rejected.

2.2 Unreal Cases

A unreal ethical optimum occurs when some ethical optimum has one or more of the variables is set arbitrarily to infinity.

> **Unreal Ethical Optimums**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N equal to infinity and/or material goods G equal to infinity, which is equal to the maximum wellbeing.

Two trivial ethical optimums exist: 1) the Utility Monster, the case of infinite material goods; 2) the Repugnant Conclusion, the case of infinite population.

2.2.1 Utility Monster

Description

The first type of unreal ethical optimum is the Utility Monster, an unreal solution in the domain of material goods.

In "Anarchy, State, and Utopia", Robert Nozick proposes the possibility that the best world is one in which a single person experiences all of the pleasure:

> Utilitarian theory is embarrassed by the possibility of utility monsters who get enormously greater sums of utility from any sacrifice of others than these others lose ... the theory seems to require that we all be sacrificed in the monster's maw, in order to increase total utility.[22]

In other words, perhaps one way to maximize all people's wellbeing is to maximize one special person's wellbeing, ignoring others.

In the statement of the Utility Monster, in the domain of material goods greater than or equal to zero, the rate of wellbeing per material goods consumed for the Utility Monster is greater than the rate of wellbeing per material goods consumed for any other person at any level of consumption. A fat, stuffed Utility Monster gains more from one more bite than a starving child gains from the same. For this to be true then, the Utility Monster's rate of wellbeing per good must increase in the entire domain of consumption.

Definition

So, we may formally define the Utility Monster:

> **Utility Monster**: For some population N_m for all states of affairs x, there is some material goods $G + 1$ with a total wellbeing U greater than the total wellbeing for some the material goods G (S.2.2.1.a).

This further implies via extrapolation that total wellbeing necessarily approaches maximum total wellbeing as the goods consumed approaches infinity. The Utility Monster's optimum is at infinite good consumed.

> **Unreal Monstrosity**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N is greater than or equal to zero and the material goods G equal to infinity, which is equal to the maximum wellbeing (S.2.2.1.b).

In other words, if the rate of the Utility Monster's wellbeing ever began to decrease with further goods consumed, the Utility Monster would not imply an ethical solution at infinity. Some other person's wellbeing would be the optimal allocation of goods.

Rejection

However, we can reject the Utility Monster as unreal (S2.2.1), and it thus can be specified away. This follows from an assumption that, as one further specification to the ethical domain conditions, we might stipulate that material infinities are excluded from ethical solutions. Because infinite goods are unobtainable in the material domain, they are not obligatory in the ethical domain:

> **Illbeing in Extremis (of goods)**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N is greater than or equal to zero and the material goods G equal to infinity, which is arbitrarily less than zero (S2.2.1c).

---

[22] Nozick, p. 41.

Assuming that positive wellbeings exist (S2.2.1d), and that thus the total wellbeing maximum must be a positive wellbeing (S2.2.1.e), and assuming a well-formed function and using the vertical line test to establish that two total wellbeings cannot correspond to one state of affairs,[23] the Utility Monster cannot both be a total wellbeing maximum and consistent with our Illbeing in Extremis condition, and thus the Utility Monster must be rejected as a total wellbeing maximum (S2.2.1.f).

2.2.2 Repugnant Conclusion

Description

The second type of unreal ethical optimum is the Repugnant Conclusion, an unreal solution in the domain of populations of persons.

In "What Should We Do About Future Generations?", Y. K. Ng elegantly puts describes the principle of mere addition as the increase of wellbeing with each subsequent additional positive-wellbeing person:

> For any given outcome A, if an alternative outcome A+ contains all the individuals in A with exactly the same utility profile as in A . . . plus a number of additional individuals all with worthwhile lives, then A + is better than or at least not worse than A.[24]

In other words, adding more population of positive wellbeing increases wellbeing overall.

By itself, the Mere Addition problem is not a problem as such, but becomes a problem without restraint. The Repugnant Conclusion is reached when the premise of "Mere Addition" is permitted to its ultimate conclusion. As Derek Parfit originally formulates it:

> For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living.[25]

In other words, one way to maximize all people's wellbeing is to maximize the number of people, but let their individual wellbeing fall to near zero.

In the statement of the Mere Addition Paradox, in the domain of population greater than zero, the rate of wellbeing per increase in population is always greater than zero. An enormous but marginally satisfied population is always better than a small but extremely ecstatic population. If this is true, because the rate of total wellbeing per population is always positive, the Mere Addition Paradox implies the Repugnant Conclusion, that the total wellbeing increases even when the state of affairs is arbitrarily overpopulated.

Definition

So, we may formally define the Repugnant Conclusion:

> **Repugnant Conclusion**: For all states of affairs x, there is some population of persons $N + 1$ with a total wellbeing U greater than the total wellbeing for some the population of persons N (S2.2.2.a).

---

[23] Note: for the above proofs, for a well-defined function there must be one and only one output for each input, thus satisfying the vertical line test ($\sim(f(x_1) \rightarrow y_{1a} \cap y_{1b} \cap y_{1a} \neq y_{1b})$).

[24] Ng, p. 237.

[25] Parfit, p. 388.

This further implies via extrapolation that total wellbeing necessarily approaches maximum total wellbeing as the population approaches infinity. The Repugnant Conclusion's optimum is at infinite population.

> **Unreal Repugnance**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N equal to infinity and material goods G greater than or equal to zero, which is equal to the maximum wellbeing (S2.2.2.b).

In other words, if any additional person decreased the wellbeing of the population, then Mere Addition would not occur indefinitely and the Repugnant Conclusion would not imply an ethical solution at infinity. Some additional person would reduce the total wellbeing below the optimal population.

Rejection

However, we can reject the Repugnant Conclusion as unreal (S2.2.2), and it can be specified away, like with the Utility Monster, via a similar exclusion conditions, this time applied to persons. Because infinite persons are unobtainable in the material domain, they are not obligatory in the ethical domain:

> **Illbeing in Extremis (of persons)**: There exists some state of affairs x where there is some total wellbeing U which is a function of the population of persons N equal to infinity and the material goods G greater than or equal to zero, which is arbitrarily less than zero (S2.2.2c).

Assuming that positive wellbeings exist (S2.2.2d), and that thus the total wellbeing maximum must be a positive wellbeing (S2.2.2.e), and again assuming a well-formed function and using the vertical line test, the Repugnant Conclusion cannot both be a total wellbeing maximum and consistent with our Illbeing in Extremis condition, and thus the Repugnant Conclusion must be rejected as a total wellbeing maximum (S2.2.2.f).

3 Ethical Zero-Value Boundary Conditions

Given the above analysis, we can arrive upon three boundary conditions upon the ethical domain in the form of zero-value points, conditions at which total wellbeing is zero: Neutrality, Satiety, and Subsistence. In all three cases, there must be some zero-value root because the Welfare Function is assumed to be representable mathematically as a continuous polynomial, and so for ranges of wellbeing that are positive ($u_i > 0$) and negative ($u_i < 0$), some wellbeing must necessarily be zero ($u_i = 0$). This follows from there being at least one case of positive wellbeing, and, for a continuous polynomial, some codomain equal to either positive or negative infinity ($u_i = \pm\infty$), and so because we have excluded unreal solutions ($+\infty$) and we only have unreal problems ($-\infty$), there must necessarily be some root.

3.1 Neutrality Condition

Given two assumptions, Ethical Problems and Incidentally Positive Ethical Domains, the existence of states of affairs with negative wellbeing (S3.1.a) and states of affairs with positive wellbeing (S3.1.b), we can assume a continuous function and apply the intermediate value

theorem (S3.1.c),[26] to conclude that there must exist some zero-value wellbeing at some value of population and goods (S3.1.d).[27]

> **Neutrality Condition**: Total wellbeing U is equal to zero at some state of affairs of neutrality with the population of persons N greater than or equal to zero and material goods G greater than or equal to zero (S3.1).

3.1 Satiety Condition

Given three assumptions, Illbeing in Extremis for positive (S3.2.a) and negative (S3.2.b) infinite material goods and Incidentally Positive Wellbeing Domains (S3.2.c), again assuming a continuous function and applying the intermediate value theorem (S3.2.d), we conclude that there must exist some zero-value wellbeing for a given population at some quantity of satiating material goods (S3.2.e).

> **Satiety Condition**: Total wellbeing U is equal to zero at some state of affairs of satiety with the population of persons N equal to one and material goods G greater than or equal to zero (S3.2).

This represents the point of satiety at which one person (a would-be Utility Monster) can no longer derive further positive wellbeing from further consumption of material goods.

3.3 Sufficiency Condition

Given three assumptions, Illbeing in Extremis for positive (S3.3.a) and negative (S3.3.b) infinite populations of persons and Incidentally Positive Wellbeing Domains (S3.3.c), again assuming a continuous function and applying the intermediate value theorem (S3.3.d), we conclude that there must exist some zero-value wellbeing for given material goods at some subsisting population of persons (S3.3.e).

> **Subsistence Condition**: Total wellbeing U is equal to zero at some state of affairs of subsistence with the population of persons N greater than or equal to zero and material goods G equal to one (S3.3).

This represents the point of subsistence at which one unit of material good can no longer produce further positive wellbeing by being distributed amongst a further population of persons (a would-be Repugnant Conclusion).

3.4 Combined Conditions

Notably, the Neutrality condition is optional for real Ethical Solutions, since it can be accommodated either by the Satiety or Subsistence conditions. Furthermore, either the Satiety or Sufficiency condition is optional if they are identical. However, the Satiety and Subsistence conditions are both necessary for real ethical solutions if there is to be Incidental Positive Wellbeing, since they must be mutually exclusive if positive wellbeing is possible. Thus, two or more zero-value wellbeing conditions are necessary in the real positive ethical domain.

Thus, starting from definitions of our unreal ethical problems, we have arrived through a deductive framework axiomatically at conditions for the ethical solutions.

---

[26] For the above proofs, all real functions are assumed mathematically continuous. Thus, for domains of values with outputs that are positive ($y_a > 0$) and negative ($y_b < 0$), some outputs must necessarily be zero ($y_{ab} = 0$).

[27] A root at zero good ($G_{izero} = 0$) would further satisfy the "neutrality of existence intuition" proposed by some authors (Broome).

4 Specifying a Real Non-Trivial Functions

From the above analysis, it can be deduced that ethical domains for real populations are possible, and expressible as functions bounded by the three zero-value conditions that avoid unreal Ethical Solutions. An ethically adequate function would have to follow at least zero-value conditions listed above. That is to say, the ethical domain could be designed according to certain design specifications, such that the unreal cases could be specified out of the domain. The specifications would at minimum omit Lossless Depopulation, the Utility Monster, and the Repugnant Conclusion. After performing such specifications, we would find a satisfactory subset of marginal welfare functions (marginal utility functions). This will give reasons to replace a conventional economic assumption, that of Diminishing Marginal Returns, with a more convoluted curvature, thus averting many of the quantitative problems that plague population ethics.

To elaborate upon the possibilities within these boundary conditions, in the following sections, we can walk through ethical functions of increasing complexity, each with a further constraint than the last, to correct for problems of triviality and unreality. In the tradition of Hurka, Ng, and Sider, we may deduce what subsets of polynomial approximation functions can characterize a real ethical domain. For this analysis, variables are defined to construct welfare function (Appendix B: Table 2) and their corresponding curves (Appendix C: Figures 1-10) to fit the zero-value conditions required.

4.1 Some Further Simplifying Assumptions

For brevity, we will make some simplifying assumptions on the functions within the bounds of the zero-value constraints provided above. (Note: these assumptions do not hold for all functions that satisfy the zero-value conditions, only special cases.)

Four mathematical assumptions are made, which have already been touched upon:

1. The functions are presented as polynomials approximations, which is reasonable to the extent that any function can be approximated by some higher order polynomial across a finite domain.[28]
2. The functions are continuous, satisfying the intermediate value theorem, which is reasonable because, even though populations of persons are discrete, excluding the possibility of fractions of persons, discrete functions can be approximated by continuous functions can approximate discrete functions, in this case by only using natural number populations.
3. The functions are well-defined, satisfying the vertical line test, which is reasonable to the extent that it is reasonable to assume that only one experiential state of wellbeing can supervene upon any one given material state of affairs, which excludes the possibility of zombies, color-inverts, and other such exotic psychological superveners.
4. The functions are extended along some real number axis between negative infinity to positive infinity, which is reasonable to the extent that the material world is finite, and

[28] Stone, M. H. (1948), "The Generalized Weierstrass Approximation Theorem", Mathematics Magazine, **21** (4): 167–184, doi:10.2307/3029750, JSTOR 3029750, MR 0027121*; **21** (5), 237–254.*

further satisfies the requirement of the polynomial approximation to be limited to a finite domain.

Three further simplifying functional assumptions are made about distributivity, a general and specific relationship between goods and populations:

> **Person-Distributed Goods**: material goods G are distributed amongst persons as a function of populations N and total material goods G_t of the form $G = f(G_t, N)$.
>
> **Equally Distributed Goods**: material goods G are *equally* distributed amongst persons as a function of populations N and total material goods G_t according to the form $G = \boldsymbol{G_t/N}$.
>
> **Completely Consumed Goods**: material goods G are *completely* by persons, such that the total material goods G_t is constant $G_t = k$.

These completionist distributive egalitarian assumptions may be considered reasonable since they are implied by the Repugnant Conclusion itself and are therefore intrinsic to its formulation.[29] But more, if the ethical solution can be shown for at least one completionist distributive egalitarian welfare function, then it follows that at least one welfare function in general satisfies the ethical solution, which is all that is meant to be shown in the first place.

Using these assumptions and the zero-value conditions, first, we will show how Linear Marginal Returns is an inadequate theory that is correctly abandoned. Secondly, we will show how the generally accepted account, Diminishing Marginal Returns, is a more satisfying theory, but one which still admits of unreal solutions. Thirdly, we will propose Inflecting Marginal Returns and show that it excludes the unreal solutions of Diminishing Marginal Returns. Fourthly, we will speculate upon the value of higher order theories of Marginal Returns (Appendix B: Table 2).

4.2 Linear Marginal Returns

Starting simply, assuming a "Linear Marginal Returns", a first order polynomial function of goods versus wellbeing, is the simplest assumption that one can make about Marginal Returns. This theory assumes that the individual wellbeing is proportional to the individual goods. As the amount of individual goods increases, the amount of individual wellbeing increases (Row 2 of "Table 4" and "Figure 1"). The individual wellbeing is positive from scarcity ($G = 0$) through infinite surplus ($G = \infty$). From intuition and experience, we know that eating more food gets us fuller than eating less food.

Fortunately "Linear Marginal Returns" has the virtue of solving the problem of Lossless Depopulation. As can be observed (in "Figure 3"), the case of zero persons is mathematically possible for a linear Marginal Returns case, but many more cases exist in the positive domain above zero, satisfying the Neutrality Condition. On the graph, at the y-axis, where the population is zero, there is a single point with total wellbeing zero. Although this case is desirable when compared with cases of negative wellbeing, it is undesirable when compared with cases of positive wellbeing, eliminating zero population as a viable solution in the ethical system.

Unfortunately, the problem resulting from the assumption of "Linear Marginal Returns" is the problem of the Utility Monster. As can be observed (in "Figure 3"), the case of one person at maximum wellbeing is mathematically possible for a linear Marginal Returns case. On the graph,

[29] Parfit, p. 388.

while population is approaching zero the individual wellbeing approaches infinity. Although a counterintuitive position, this case makes infinite personal greed a potential solution in the ethical system.

Generally, as can be observed (in "Figure 3"), under linear Marginal Returns, there is no real relative maximum reached by the total wellbeing function. The graph instead consists of a continuous flat line, across the positive range of populations, of constant total wellbeing across the domain. The function is static, always at a "best" case scenario, but never suggesting one scenario over another, nothing that is finite, non-trivial, and unique.

Because "Linear Marginal Returns" incurs these problems, it has been correctly abandoned as a viable ethical principle and replaced by the more robust Diminishing Marginal Returns.

4.3 Diminishing Marginal Returns

Assuming a "Diminishing Marginal Returns", a second order function of goods versus wellbeing, is the most common assumption that one can make about Marginal Returns, long assumed to be an a priori axiom of human behavior.[30] This theory assumes that the individual wellbeing is related to the individual goods parabolically, following a second-order concave-down polynomial. As the amount of individual goods increases, the amount of individual wellbeing increases, until a point of satiety, after which point, as the amount of individual goods increases, the amount of individual wellbeing decreases (Row 3 of "Table 4" and "Figure 2"). As the graph shows, the individual wellbeing is positive from scarcity ($G = 0$) through to maximum goods consumable before suffering ($G = G_{sat}$). From intuition and experience, we know that eating more food is good, until we are stuffed, at which point, eating more food is bad.

Fortunately, besides matching our experiences, "Diminishing Marginal Returns" has the virtue of solving the problem of the Utility Monster, satisfying the Satiety Condition. As can be observed (in "Figure 4"), the case of one person is mathematically possible for a diminishing Marginal Returns case, but many more cases exist in the positive domain with populations of more than one person. In other words, because wellbeing diminishes across a range of consumption, individual wellbeing is capped at a maximum. Thus, no individual person can ever reach infinity wellbeing, and thus no singular person can ever become a wellbeing monster.

Unfortunately, this common assumption also results in problems.

1) The primary problem resulting from the assumption of "Diminishing Marginal Returns" is the problem of Mere Addition. As can be observed (in "Figure 4"), the addition of more and more positive persons results in the addition of more and more positive wellbeing, in the case of diminishing Marginal Returns. The graph increases in total wellbeing, across the positive range of populations, never returning back towards zero.
2) The secondary problem is that of the Repugnant Conclusion. As can be observed, under Diminishing Marginal Returns, there is no relative maximum reached by the total wellbeing function, under the diminishing Marginal Returns case. The function achieves its absolute maximum total wellbeing only far off the graph approaching infinity. The function drifts off indefinitely, never actually achieving a "best" case scenario that is finite, non-trivial, and unique.

[30] Mises, Ludwig Von. *Human Action*. The Ludwig von Mises Institute, (1998), p. 119.

Because “Diminishing Marginal Returns” incurs these problems, it must be abandoned as a viable ethical principle. Most economic theory assumes a function of “Diminishing Marginal Returns”, but to refine ethical theories and avoid corresponding ethical conundrums, higher-order Marginal Returns might better serve. As it turns out, a second order equation does not necessarily have enough design specifications to specify out three or more boundary conditions.

4.4 Delayed Diminishing Marginal Returns

Assuming a “Delayed Diminishing Marginal Returns”, identical to regular diminishing Marginal Returns, but translated positively on the x-axis such that both x-intercepts are greater than zero, is another possible assumption. As the amount of individual goods increases, the amount of individual wellbeing increases, until a point of subsistence, after which individual wellbeing increases more, until a point of satiety, after which point, as the amount of individual goods increases, the amount of individual wellbeing decreases. From intuition and experience, we know that eating below some certain quantity food is bad, because we starve, and that only eating more than that quantity is a net good, after staving off hunger. Delayed Diminishing Marginal Returns is consistent with Parfit’s own observations of the pitfalls of Diminishing Marginal Returns.[31] Indeed, this is an improved assumption for two reasons: firstly, its diminishing trend still avoids the wellbeing monster by maxing out individual wellbeing; secondly, it’s delayed positivity avoids the repugnant conclusion by minning out individual wellbeing. In other words, Delayed Marginal Returns satisfies the Satiety Condition and the Subsistence Condition. Unreal zeros and infinities have been specified out of the formulation.

4.5 Inflecting Marginal Returns

Assuming a “Inflecting Marginal Returns”, a third order function of goods versus wellbeing, is perhaps a more promising assumption that one can make about Marginal Returns. This theory assumes that the individual wellbeing is related to the individual goods via a third order, or possibly higher order, concave-down polynomial. The function has three extrema: one relative maximum, one relative minimum, and one flat inflection point. The flat inflection point exists at the case of zero goods. The relative minimum and relative maxima each exist at the cases of some positive quantity of goods, at a point of return on investment and at a point of satiety respectively. As the amount of individual goods increases, the amount of individual wellbeing decreases, until a point of return on investment, after which point, as the amount of individual goods increases, the amount of individual wellbeing increases, until the point of satiety, after which point, as the amount of individual goods increases, the amount of individual wellbeing decreases, (Row 4 of “Table 4”, “Figure 5”, and “Figure 7”). As the graph shows, the individual wellbeing is negative from scarcity ($G = 0$) through to subsistence, the minimum goods consumable after suffering ($G = G_{sub}$); then the individual wellbeing is positive through to fullness, the maximum goods consumable before suffering ($G = G_{sat}$). From intuition and experience, we know that eating more food is good, until we are stuffed, at which point, eating more food is bad; but we also know that working for food is bad, until a return on investment, at which point, working for food becomes good. In other words, it satisfies the Neutrality, Satiety, and Subsistence Conditions. Like Delayed Diminishing Marginal Returns, the assumption of

[31] “On some versions of the *Law of Diminishing Marginal Returns*, this is just what is implied. On these versions, each unit of resources produces more utility if it is given to people who are worse off, so that the most productive distribution will be the one where everyone's life is barely worth living. There is here an obvious oversight. Large amounts of resources are needed to make each person's life even reach the level where life begins to be worth living.” (Parfit 525)

Inflecting Marginal Returns does not result in any of the problems of trivial and unreal wellbeing—monstrosity and repugnance—listed above ("Table 5").

4.6 Fourth Order Marginal Returns (and further?)

Although the third order function has resolved our three problems, further problems may be suspected to emerge due to the behavior of the specified domain. For example, for the third order curve, in regions of increasingly negative goods, wellbeings are increasingly positive, a clearly incoherent result. Negative goods are probably excluded from the range of physical possibility anyway, rendering the problem more theoretical than practical. But, even if negative goods are materially possible—whether via some conceptualization of indebted value or antimatter particles, we shall not speculate here—intuition protests that negative goods surely cannot produce positive value. Fortunately, this problem is as easily resolved as it is stated. By changing the third order function to a fourth order function, the graph becomes globally concave down, meaning that negative goods correspond to negative wellbeing, while no other attributes of the curve change except the negative region of goods (Row 5 of "Table 4", "Figure 6", and "Figure 8"). (Note: although we have eliminated positive infinities, we are now left with negative infinities, which might be objected to as well, if humans are finite beings. This would suggest that some function other than a polynomial would more accurately describe the behavior of Marginal Returns, which is beyond our present scope.) As further such problems arise, the polynomial can be altered to accommodate them, each problem corresponding to a design specification corresponding to an addition order to the polynomial.

5 Discussion, Objections, Responses

So, rather than accepting unreal ethical solutions, we have shown that treating them as boundary conditions to be excluded with design specifications on a welfare function, which has intuitive results. Any satisfactory ethical domain would have to be specified to avoid these problems, and thus any ethical domain with these problems would be underspecified to the extent that it admits of impossible solutions. The result might offer a solution that we might dub a "Resplendent Conclusion" in contrast to Parfit's "Repugnant Conclusion":

> **A Resplendent Conclusion:** There exists some state of affairs x in the real domain of populations of persons N and the material goods G at which total wellbeing U is maximized.

The Resplendent Conclusion might be a fair candidate for a Theory X.

However, even after all this, admittedly, this result might have its own substantive problems, not resolving Repugnance and Monstrosity but merely replacing them with other issues. So, the next steps are to consider these new problems and attempt to accept them as intuitive or resolve them with further constraints:

Dealing with the Fertile/Sterile Conclusions

Because there is an ethical solution at some discrete, finite population, there are real ethical domains in which is becomes preferrable to marginally increase or decrease the population, implying counterintuitive conclusions respectively:

> **A Fertile Conclusion:** For some states of affairs x of small populations of persons N of reasonable wellbeing U, there are be some states of affairs x_{max} with some larger population of persons N_{max} with even greater wellbeing u_{max}.

And:

> **A Sterile Conclusion:** For some states of affairs x of large populations of persons N of reasonable wellbeing U, there are some states of affairs x_{max} smaller population of persons N_{max} with even greater wellbeing u_{max}.

In other words, there might be some optimal possible population, but if there is, then necessarily all larger and smaller populations are sub-optimal. The population control techniques of breeding and culling become ethical imperatives. The Resplendent Conclusion implies the Fertile/Sterile Conclusion, so we cannot accept the solution without these problems as well. The Fertile/Sterile Conclusions are at odds with intuitions of reproductive freedom, and they may even imply some level of breeding a population, via forcible fertilization/sterilization, may be justified by larger prospective wellbeings. Indeed, the Sterile Conclusion may even imply that some level of culling a population can be justified.

Thus, in some sense, the coupling of the Resplendent result with the Fertile/Sterile Conclusions reinforces the supposition that there is no solution to population ethics (Arrhenius, 1999). But, in another sense, the Fertile/Sterile Conclusions are conceptually more sensible than the Repugnant Conclusion, because the Fertile/Sterile Conclusions present real, finite, non-trivial, and unique solutions to problems of welfare. Thus, although unattractive, the Fertile/Sterile Conclusions are real, unlike the Repugnant Conclusion, which is both unattractive and unreal.

Furthermore, there are at least three straightforward ways to dampen the counter-intuitiveness of the Fertile and Sterile Conclusions.

First, biologically, in terms of scarcity and plenitude, the theory presented here does not account for the effects of poor wellbeing on the medical conditions of persons. In reality, persons who spend a lot of time in material conditions that maintain high wellbeing have greater health generally and thus tend to persist, and are better able to live longer lives and reproduce more; while persons who spend a lot of time in material conditions that maintain low wellbeing conditions have lower health generally and thus tend to perish, and are wont to live shorter lives and reproduce less. Thus, in a sense, population ethics has a self-correcting mechanism built-in to the material ecologies of creatures, and thus does not need creatures to feel these intuitions, the Fertile and Sterile Conclusions, because normative intuition for either conclusion is redundant to natural process. In other words, animals do not need an ethical reason to give birth or die off, their bodies just do it for them according to biological principles. For example, in the surplus human economy of modern Earth, we rarely need to breed or cull ourselves, because we almost always have enough; in the shortage economy of a Martian colony, we may not need to breed or cull ourselves either, because we may simply find ourselves in conditions in which persons are birthing due to resource surplus and dying off due to resource shortages without any human intervention. Thus, in both cases, if ecology takes care of itself, then in neither case do the human agents have to act upon an intuition towards breeding or culling, so the Fertile and Sterile Conclusions are largely redundant.

Second, epistemically, in terms of certainties and uncertainties, the theory presented here should not be read as implying that we are confident about our welfare conditions in any given case.

Indeed, we are almost never in an epistemic situation in which we can judge how additional births or deaths will affect overall wellbeing, and we are thus almost never in the position of ethical authority to oblige or permit addition and subtraction of population, so the Fertile and Sterile Conclusions, while perhaps true, are perhaps never actionable. Indeed, although the mathematics we have used here is suggestive of precision, this is merely the precision of theory; in the messiness of reality, these welfare functions have wide error bars, and thus more gray area than black and white area, and thus are not to be mistaken for definitive policy prescriptions. The only cases approaching certainty would be cases well below optimum welfare (cases with miniscule populations in situations of wasted plenty), or well above optimum welfare (cases with massive populations in situations of immiserating scarcity), assuming such conditions can be assessed demographically. In other words, we can only make the choice to birth or die under when we are unrealistically confident in our choice. For example, because we live in an enormous, uncontrolled environment on Earth, the welfare situation is often unclear, and we lack the certainty necessary to make population ethics actionable; in contrast, if we lived in a small, controlled environment like a Martian colony, the situation would have clear boundary conditions, minimum and maximum viable resource and population sizes, and we would have more certainty about our population ethics. Thus, in the latter case, but not the former case, we would feel stronger certainty in implementing the Fertile and Sterile Conclusions because we could actually make relatively precise calculations about the upper and lower bounds of good outcomes.

Third, normatively, in terms of permissions and obligations, the theory presented here is agnostic about whether betterness is obligatory, permissible, pro tanto, or some other conditional consideration. If betterness is obligatory, then the Fertile and Sterile Conclusions follow as obligatory, which seems quite strict, and is perhaps is prima facie reason to reject the Resplendent Conclusion. If betterness is merely permissible, then the Fertile and Sterile Conclusions follow as merely permissible, which seems less strict but perhaps too loose. If betterness is pro tanto overridden by other duties, then the Fertile and Sterile Conclusions only follow if not overridden (e.g., other personal goods, other public goods, etc.). If betterness is not obligatory/permissible but maintaining a positive wellbeing is obligatory/permissible, then the Fertile and Sterile Conclusions follow only when conditions drop below the positive level: in other words, the worst case scenario would be the euthanasia of miserable. Assuming that population ethics is only conditionally obligatory/permissible in this last manner, it would be no wonder that we do not feel the pull of the Fertile and Sterile Conclusions as intuitions, because we rarely find ourselves in such conditions. Indeed, from this hint we may further suspect that our intuitions against the Fertile and Sterile Conclusions may only be a feature of a surplus world, and we may come to acquire the reverse intuitions in a world of shortage. For example, because we live in the resource rich environment of Earth, where we enjoy incomplete consumption, we rarely find it necessary to breed or cull ourselves; in contrast, if we lived in a tight and resource constrained environment, suffering from the triage conditions of complete consumption, like in a Martian colony, we would surely find it necessary in certain cases to breed or cull ourselves. Thus, in the latter case, but not the former case, we would feel stronger intuitive support for the Fertile and Sterile Conclusions.

Such are some considerations that make the Fertile/Sterile Conclusions plausible. And, although we think these above explanations will satisfy upon reflection, more detailed work resolving the Fertile/Sterile Conclusions should be the subject of future ethical endeavor. The open question

here offers further philosophical questions about the subject of population ethics and an invitation to expanded inquiry.

Dealing with the Egoism/Universalism Disagreement

Also, there is an ethical solution for optimizing the total wellbeing of a population (the universalist priority), and there is a different ethical solution for optimizing the individual wellbeing of a person (the egoistic priority):

> **The Egoism/Universalism Disagreement:** For states of affairs x of populations of persons N of wellbeing U, the states of affairs $x_{i,max}$ providing the optimum individual wellbeing of any given person $u_{i,max}$ will be different from the states of affairs $x_{N,max}$ providing optimum total wellbeing of the population $U_{N,max}$, unless trivially the environment only materially supports a single person.

In other words, the best for the person is not the best for the population. So, the marginal increase of the wellbeing of the population may incidentally (and will eventually) result in the marginal decrease of the wellbeing of the average person. The Resplendent Conclusion implies the Egoism/Universalism Disagreement, to the extent that optimizing requires differential calculus, and the derivative of a function is not necessarily mathematically the same as the derivative of that function times an independent variable. The Egoism/Universalism Disagreement is at odds with intuitions that the interests of communities are aligned with the interests of individuals, and even suggests that the two are necessarily incidentally at odds, a suggestion consistent with Henry Sidgwick's Dualism of Practical Reason.

The coupling of the Resplendent result with the Egoism/Universalism Disagreement reinforces the supposition that egoism and universalism prescribe different outcomes, but this account is ultimately intuitive, for several reasons, and thus is easily accepted.

First, we would expect egoism to differ from universalism, because we frequently encounter, as individuals, sacrifices that we make for the communal good. Indeed, in some sense, this is different version of the submission of certain privileges to the Leviathan in the Hobbesian social contract. And, considered in terms of reproductive ethics, we can observe the simple dynamic that, in any given family, order to achieve the best outcome for the family, parents must give away some of their individual wellbeing to the children.

Second, the difference between egoism and universalism is not repugnantly large, and so the individual and community are not in complete conflict, only conflict on margin near their respective optimums. Indeed, according to the Resplendent Conclusion, some level of individual wellbeing is required for total wellbeing, which is much better than the Repugnant Conclusion, where individual wellbeing must be minimized for the total wellbeing. Indeed, per the Resplendent Conclusion the individual need not sacrifice much at all from their maximum to achieve the total maximum. The individual must have some children, but not too many, and must sacrifice some of their own pleasure, but not too much.

Third, the difference between egoism and universalism has a positive construal as the basis for establishing a range reasonably acceptable behavior, as well as a basis for describing some actions as supererogatory. Egoists can be dubbed as behaving rationally, on the low end of a egoistic/universalistic range; and universalists can be also dubbed as behaving rationally, on the high end of the range; and we need to suggest that everyone must rationally follow the same priorities. Furthermore, to the extent that there is a difference, the discrepancy between selfless

and selfish optimums differentiates heroics as distinct from hedonics, sacrifice as distinct from satisfaction, giving as distinct from greed.

Such are some considerations that make the Egoism plausible.

Conclusion

Regardless, relinquishing simple assumptions about Diminishing Marginal Returns and partaking instead of more complex curvature offers a solution to long-standing problems in population ethics, satisfying a set of zero-value conditions that exclude trivial and unreal ethical solutions. Furthermore, this does not require the introduction of new values or the destruction of old ones, rather it remains consistent with Parfit's Beneficence Principle. Perhaps the Inflecting domain opens new ethical problems that the Diminishing domain foreclosed. But, because of its finely tunable parameters and its descriptive elegance, these new assumptions offers a promising replacement for traditional ethical and economical considerations. The curve potentially resolves Parfit's ethical problem The Repugnant Conclusion, amongst others, and offers an avenue towards his unified ethical theory: "Theory X".

Acknowledgements

Special thanks to Jacob Mills and Daniel Flores of Houston Community College for their feedback on early drafts, to the planning committee of the 2018 HCC Humanities and Philosophy for the opportunity to present this project, and to David Phillips and colleagues at the University of Houston for many metaethical discussions.

Appendix A: Proofs

(Notation: U Total wellbeing, u_i Individual wellbeing, u_m Marginal wellbeing, N Population, G Individual Goods, G_t Total Goods, x states of affairs.)

Table 1: Formal Proofs of the Boundary Conditions of the Ethical Domain

Step		Formal Description	Informal Description
S1.1		$\forall \boldsymbol{x}[\boldsymbol{U}(\boldsymbol{N}(\boldsymbol{x}) \geq \boldsymbol{0})(\boldsymbol{G}(\boldsymbol{x}) \geq \boldsymbol{0}) = \boldsymbol{U_x}]$	Ethical Domain
S1.2		$\exists \boldsymbol{x}[\boldsymbol{U}(\boldsymbol{N}(\boldsymbol{x}) \geq \boldsymbol{0})(\boldsymbol{G}(\boldsymbol{x}) \geq \boldsymbol{0}) < \boldsymbol{0}]$	Ethical Problems
S1.3		$\exists \boldsymbol{x}[\boldsymbol{U}(\boldsymbol{N}(\boldsymbol{x}) \geq \boldsymbol{0})(\boldsymbol{G}(\boldsymbol{x}) \geq \boldsymbol{0}) > \boldsymbol{0}]$	Ethical Solutions
S1.4		$\exists \boldsymbol{x}[\boldsymbol{U}(\boldsymbol{N}(\boldsymbol{x}) \geq \boldsymbol{0})(\boldsymbol{G}(\boldsymbol{x}) \geq \boldsymbol{0}) = \boldsymbol{U_{max}}]$	Ethical Optimums
S2.1.1	*show*	$\sim\exists x[U(N(x) = 0)(G(x) \geq 0) = U_{max}]$	No Lossless Depopulation
S2.1.1.a		$\exists x[U(N(x) = 0)(G(x) \geq 0) = U_{max}]$	Lossless Depopulation
S2.1.1.b		$\exists x\left[U(N(x) = 0)(G(x) \geq 0) = \frac{0}{0}\ or\ 0\right]$	Indeterminate/Zero Nonbeing
S2.1.1.c		$\forall x\left[U(N(x) \geq 0)(G(x) \geq 0) \leq \frac{0}{0}\ or\ 0\right]$	Via definition of maximum
S2.1.1.d		$\exists x\left[U(N(x) \geq 0)(G(x) \geq 0) > \frac{0}{0}\ and\ 0\right]$	Incidentally Non-negative/ Determinate wellbeing
S2.1.2	*show*	$\sim\exists x[U(N(x) \geq 0)(G(x) = 0) = U_{max}]$	No Lossless Deprivation
S2.1.2.a		$\exists x[U(N(x) \geq 0)(G(x) = 0) = U_{max}]$	Lossless Deprivation
S2.1.2.b		$\forall x[(G(x) = 0) \rightarrow (N(x) = 0)]$	Material Personhood
S2.1.2.c		$\exists x[U(N(x) = 0)(G(x) = 0) = U_{max}]$	Special Lossless Depopulation
S2.2.1	*show*	$\sim\forall x[U(N_m)\big(G(x)\big) < U(N_m)(G(x) + 1)]$	No Utility Monster
S2.2.1.a		$\forall x[U(N_m)\big(G(x)\big) < U(N_m)(G(x) + 1)]$	Utility Monster
S2.2.1.b		$\exists x[U(N(x) \geq 0)(G(x) = +\infty) = U_{max}]$	Via extrapolation
S2.2.1.c		$\exists x[U(N(x) \geq 0)(G(x) = +\infty) \ll 0]$	Illbeing in extremis (for goods)
S2.2.1.d		$\exists x[U(N(x) \geq 0)(G(x) \geq 0) > 0]$	Positive wellbeing exists.
S2.2.1.e		$\exists x[U_{max} > 0]$	Via definition of maximum
S2.2.1.f		$\sim\exists x[U(N(x) \geq 0)(G(x) = +\infty) = U_{max}]$	Via vertical line test for well-formed functions
S2.2.2	*show*	$\sim\forall x[U(G_r)\big(N(x)\big) < U(G_r)(N(x) + 1)]$	No Repugnant Conclusion
S2.2.2.a		$\forall x[U(G_r)\big(N(x)\big) < U(G_r)(N(x) + 1)]$	Repugnant Conclusion

S2.2.2.b		$\exists x[U(N(x) = +\infty)(G(x) \geq 0) = U_{max}]$	Via extrapolation
S2.2.2.c		$\exists x[U(N(x) = +\infty)(G(x) \geq 0) \ll 0]$	Illbeing in extremis (for persons)
S2.2.2.d		$\exists x[U(N(x) \geq 0)(G(x) \geq 0) > 0]$	Non-negative wellbeing exists.
S2.2.2.e		$\exists x[U_{max} > 0]$	Via definition of maximum
S2.2.2.f		$\sim\exists x[U(N(x) = +\infty)(G(x) \geq 0) = U]$	Via vertical line test for well-formed functions
S3.1	*show*	$U(N_{neut} \geq 0)(G_{neut} \geq 0) = 0$	Root 1: Neutrality.
S3.1.a		$\exists x[U(N(x) \geq 0)(G(x) \geq 0) < 0]$	Negative wellbeing exists.
S3.1.b		$\exists x[U(N(x) \geq 0)(G(x) \geq 0) > 0]$	Positive wellbeing exists.
S3.1.c		$\exists x[U(N(x) \geq 0)(G(x) \geq 0) = 0]$	Via intermediate value theorem
S3.1.d		$U(N_{neut} \geq 0)(G_{neut} \geq 0) = 0$	Existential Instantiation
S3.2	*show*	$U(N \geq 0)(G = G_{sat}) = 0$	Root 2: Satiety
S3.2.a		$\exists x[U(N(x) \geq 0)(G(x) = +\infty) \ll 0]$	Illbeing in Extremis (for goods)
S3.2.b		$\exists x[U(N(x) \geq 0)(G(x) = -\infty) \ll 0]$	Illbeing in Extremis (for goods)
S3.2.c		$\exists x[U(N(x) \geq 0)(+\infty > G(x) > -\infty) > 0]$	Positive wellbeing exists.
S3.2.d		$\exists x[U(N(x) \geq 0)(G(x) = G_{sat}(x)) = 0]$	Via intermediate value theorem
S3.2.e		$U(N \geq 0)(G = G_{sat}) = 0$	Existential Instantiation
S3.3	*show*	$U(N = N_{sub})(G \geq 0) = 0$	Root 3: Subsistence
S3.3.a		$\exists x[U(N(x) = +\infty)(G(x) \geq 0) \ll 0]$	Illbeing in Extremis (for persons)
S3.3.b		$\exists x[U(N(x) = -\infty)(G(x) \geq 0) \ll 0]$	Illbeing in Extremis (for persons)
S3.3.c		$\exists x[U(+\infty > N(x) > -\infty)(G(x) \geq 0) > 0]$	Positive wellbeing exists.
S3.3.d		$\exists x[U(N(x) = N_{sub}(x))(G(x) \geq 0) = 0]$	Via intermediate value theorem
S3.3.e		$U(N = N_{sub})(G \geq 0) = 0$	Existential Instantiation

Appendix B: Tables[32]

Table 2: Order of Welfare Function

Order	Marginal Returns	Individual Welfare Functions	Total Egalitarian Welfare Functions
0th	None[1]	u_i	$U = u_i N^{*\circ}$
1st	Linear	$u_i = G - G_{neut}$	$U = N\left(\frac{G_t}{N} - G_{neut}\right)$
2nd	Diminishing	$u_i = (G)(G_{sat} - G)$	$U = \left(\frac{G_t}{N}\right)\left(G_{sat} - \frac{G_t}{N}\right)N$
2nd	Delayed Diminishing**	$u_i = (G - G_{neut})(G_{sat} - G)$ $G_{neut} \neq 0$	$U = N\left(\frac{G_t}{N} - G_{neut}\right)\left(G_{sat} - \frac{G_t}{N}\right)$ $G_{neut} \neq 0$
3rd	3rd Order Inflecting**	$u_i = (G - G_{neut})(G_{sat} - G)(G - G_{sub})$	$U = N\left(\frac{G_t}{N} - G_{neut}\right)\left(G_{sat} - \frac{G_t}{N}\right)\left(\frac{G_t}{N} - G_{sub}\right)$
4th	4th Order Inflecting**	$u_i = (G)(G - G_{neut})(G_{sat} - G)(G - G_{sub})$	$U = N\left(\frac{G_t}{N}\right)\left(\frac{G_t}{N} - G_{neut}\right)\left(G_{sat} - \frac{G_t}{N}\right)\left(\frac{G_t}{N} - G_{sub}\right)$
Note*	The basic function of the utilitarian calculus simplified for a perfectly egalitarian case.[33]		
Note°	While the Completely Consumed Goods assumption is true in super-saturated populations ($N > N_{i,opt}$), where not every person can achieve optimum wellbeing, it is surely not necessarily true in sub-saturated populations ($N < N_{i,opt}$), where every person can achieve optimum wellbeing and should not consume more lest their wellbeing diminish. Thus, the above functions should be treated as more precise in the former case but less precise in the latter case. In order to capture sub-saturated populations, we can approximate that individuals tend to consume up to and not past their individual optimum wellbeing ($u_{i,opt}$), and thus can assume a linear relationship between population and total wellbeing ($U = Nu_{i,opt}$), reflected in the "sub"-saturated functions in the figures below.		

[32] For the figures included, the neutrality, subsistence, satiety and total goods were arbitrarily assumed to be $G_{ineut} = 0$, $G_{isub} = 1$, $G_{isat} = 5$, and $G_t = 15$ respectively with a scalar factor adjustment to force consistently $u_{i,opt} = 6.25$.

[33] Bentham, p. 7.

Appendix C: Figures

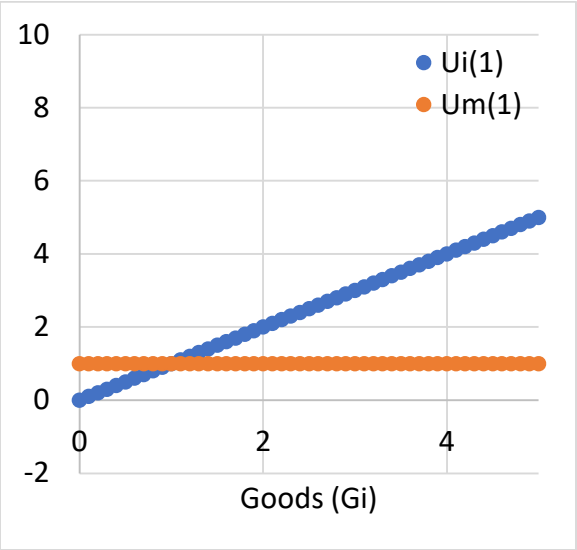

Figure 1:Linear Marginal Returns

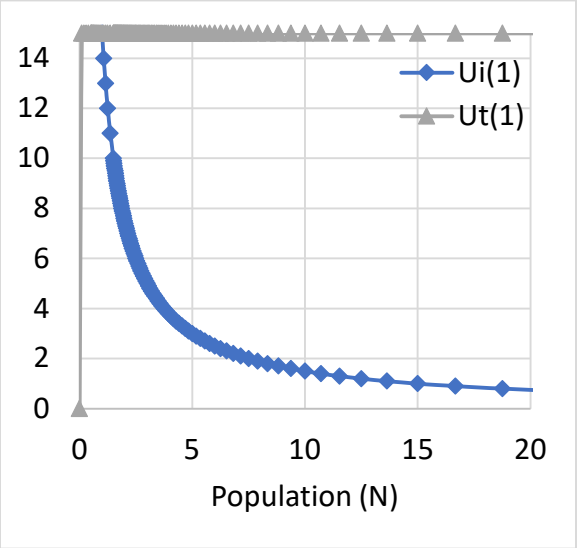

Figure 6:Linear Total Welfare

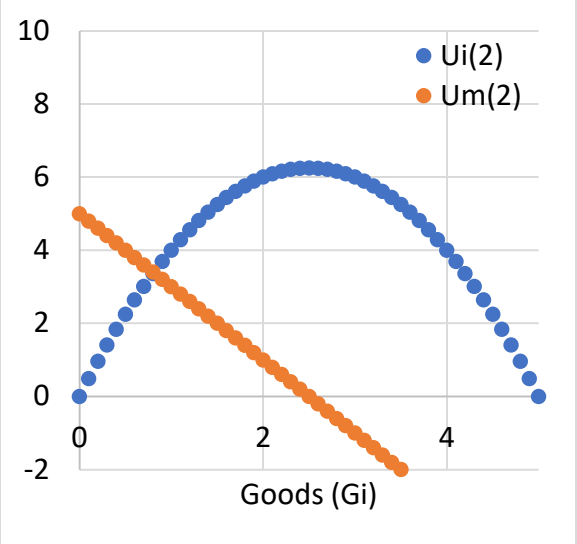

Figure 2:Diminishing Marginal Returns

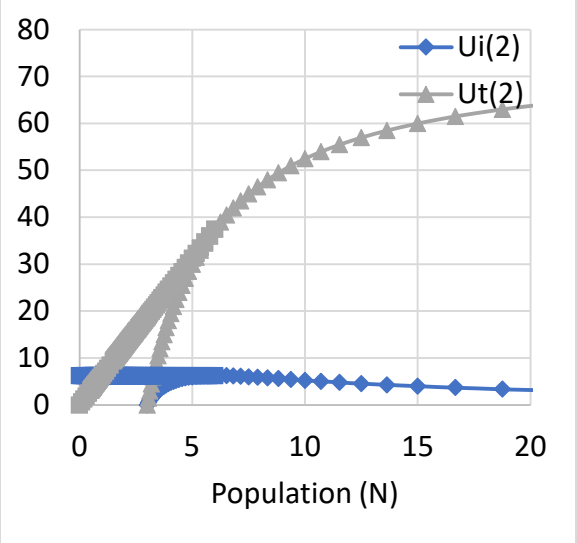

Figure 7:Diminishing Total Welfare

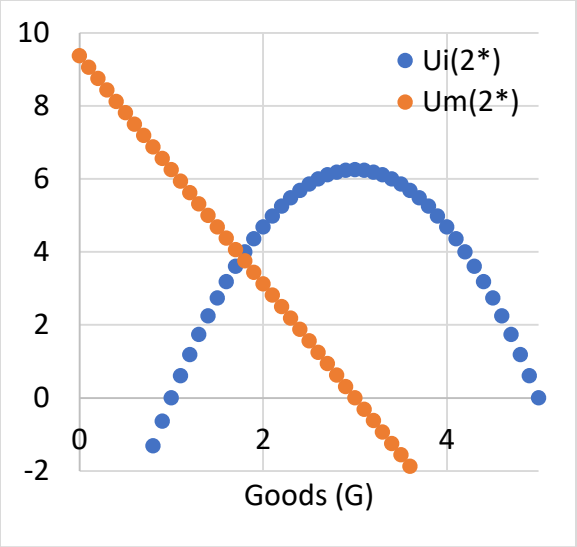

Figure 3: Delayed Diminishing Marginal Returns

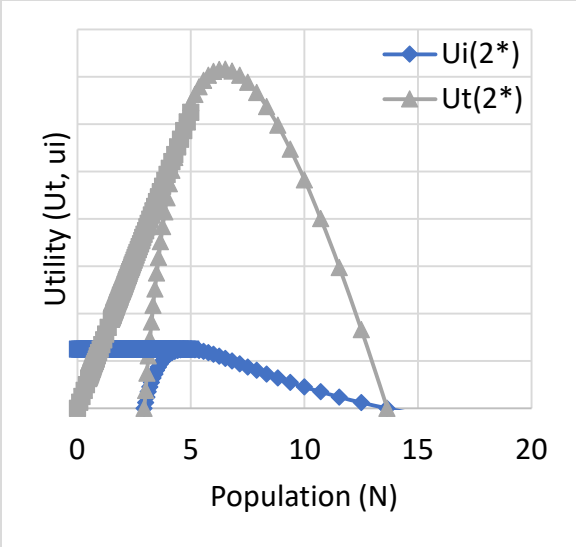

Figure 8: Delayed Diminishing Total Welfare

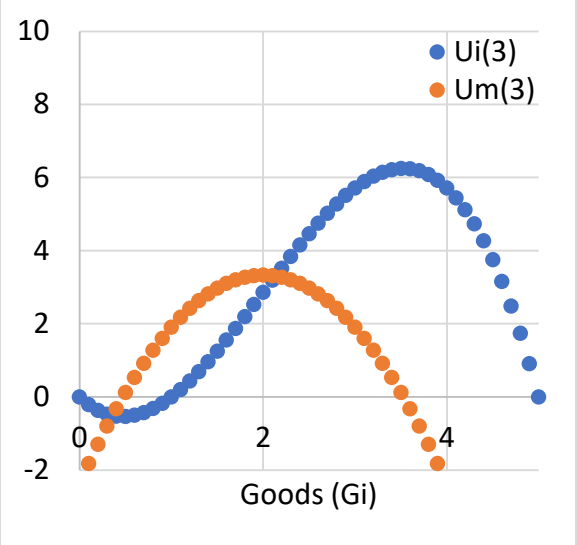

Figure 4:Inflecting Marginal Returns (3rd Order)

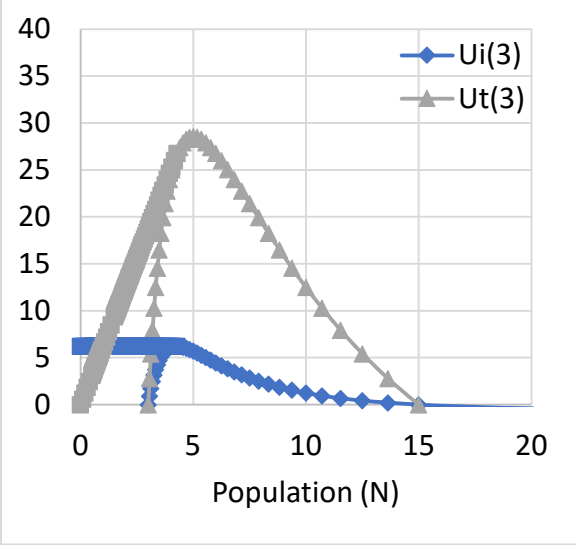

Figure 9: Inflecting Total Welfare (3rd Order)

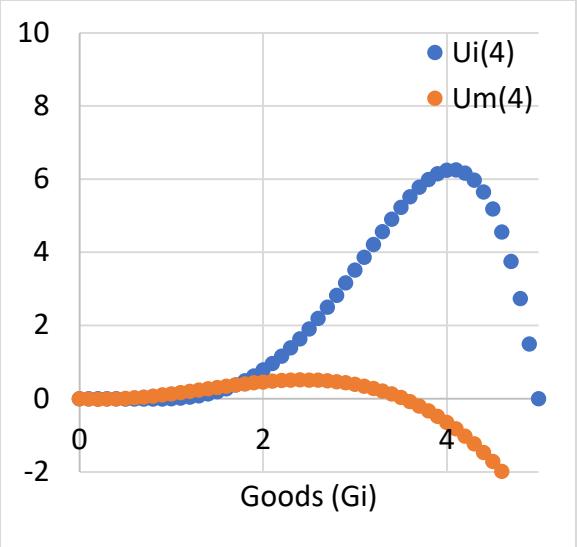

Figure 5: Inflecting Marginal Returns (4rth Order)

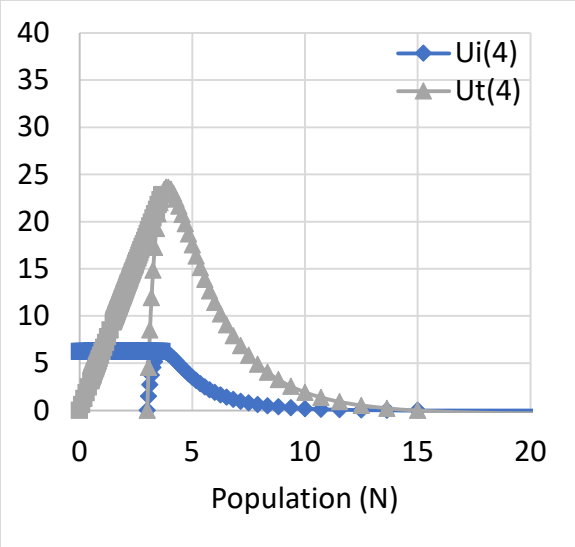

Figure 10:Inflecting Total Welfare (4rth Order)