

Ponthiere, Gregory

Working Paper

Are Long-Lived Persons Utility Monsters?

GLO Discussion Paper, No. 1186

Provided in Cooperation with:
Global Labor Organization (GLO)

Suggested Citation: Ponthiere, Gregory (2022) : Are Long-Lived Persons Utility Monsters?, GLO Discussion Paper, No. 1186, Global Labor Organization (GLO), Essen

This Version is available at:
<http://hdl.handle.net/10419/265369>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Are Long-Lived Persons Utility Monsters?

Gregory Ponthiere¹

Abstract:

Nozick's "utility monster" – a being who is more efficient than other persons at transforming resources into well-being – is often regarded as deeply impossible, on the ground of the incapacity of a single person to have a life that is better than a large number of other lives. In this article, I defend a purely marginalist view of the "utility monster", that is, that the primary characteristic of a "utility monster" is a higher sensitivity, at the margin, of well-being to resources, rather than a larger total well-being. I propose three purely marginalist accounts of "utility monster" and I introduce the related concept of "collective utility monster", in order to account for the collective predation of (almost) all resources by a group of persons. I argue that, although a long-lived person, if taken separately, could hardly belong to the category of "utility monster", a large group of long-lived persons can, under some conditions, belong to the category of "collective utility monster". In the light of the increasingly large proportion of cohorts reaching very old ages nowadays, Nozick's objection against utilitarianism turns out, after a thorough review, to be most relevant for real-world aging societies.

Keywords: longevity, mortality, inequalities, utilitarianism, Nozick's utility monster.

JEL codes: I31, J10, J18.

¹ UCLouvain, Hoover Chair in Economic and Social Ethics, and Global Labour Organization (GLO). Address: Collège Dupriez, 3 place Montesquieu, 1348 Louvain-la-Neuve, Belgium. E-mail: gregory.ponthiere@uclouvain.be. See also <https://glabor.org/user/gregoryponthiere/>

1. Introduction

In his critique of utilitarianism, Nozick (1974, p. 41) argued:

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. For, unacceptably, the theory seems to require that we all be sacrificed in the monster's maw, in order to increase total utility.

The extent to which the utility monster case can serve as a decisive argument against utilitarianism has been widely debated among philosophers. In *Reasons and Persons*, Parfit (1984, p. 389) questioned the coherence of Nozick's utility monster scenario:

As described by Nozick, such a person is a deep impossibility. The world's population is now several billion. Let us imagine the wretchedness of all these people if they are denied anything above starvation rations, and all other resources go to Nozick's imagined Monster. Nozick tells us to suppose that this imagined person would be so happy, or have a life of such high quality, that this is the distribution that produces the greatest sum of happiness, or the greatest amount of whatever makes life worth living. How can this be true, given the billions left in wretchedness that could be so easily relieved by a small fraction of this Monster's vast resources? For this to be true, this Monster's quality of life must be *millions* of times as high as that of anyone we know. Can we imagine this? Think of the life of the luckiest person that you know, and ask what a life would have to be like in order to be a million times as much worth living. [...] It seems a fair reply that we cannot imagine, even in the dimmest way, the life of this Utility Monster. And this casts doubt on the force of the example.

Parfit argues that Nozick's "utility monster" is deeply impossible, and, hence, that Nozick's "utility monster" example cannot serve as a convincing criticism against utilitarianism.

However, when examining the Repugnant Conclusion, Parfit underlines that the imagined population Z – where billions of persons live a very poor life – can be interpreted as another – more plausible – “utility monster”. Parfit (1984, p. 389) underlines that:

This imagined population is another Utility Monster. The difference is that the greater sum of happiness comes from a vast increase, not in the quality of one person’s life, but in the number of lives lived. And *my* Utility Monster is neither deeply impossible, nor something that we cannot imagine.

Parfit thinks that a possible kind of “utility monster” could take the form of an extremely large population, favouring, from a utilitarian perspective, the sacrifice of each of its members in terms of quality of lives for the sake of increasing the quantity of lives.²

In this article, I will re-examine the conditions of *existence* of Nozick’s “utility monsters”. I will focus on the particular point raised by Parfit (1984) concerning the implausibility of “utility monsters”, and leave other issues aside.³ In order to re-examine whether or not the existence of Nozick’s “utility monsters” is as deeply impossible as Parfit argued, I will consider particular candidates for “utility monsters”, who are neither imagined beings with an extraordinary capacity to enjoy their life, nor imagined extremely large populations (as population Z), but (groups of) persons who have a *longer life* than other persons. This article will examine the formal relationship between inequalities in longevity and the existence of Nozick’s “utility monsters” in a general class of resource distribution problems.

At this stage, it should be stressed that our study of the relationship between “utility monsters” and long-lived persons takes longevity inequalities *as given*, and, as such, differs

² More recently, Briggs and Nolan (2015) examined the relation between Nozick’s monster and the fission of one person into multiple persons, and deduced new versions of Parfit’s Repugnant Conclusion.

³ Parfit’s criticism of Nozick’s utility monsters as deeply impossible is not the unique way of questioning Nozick’s argument. In a recent article, Chappell (2021) argued that the utility monster criticism loses its intuitive power in the case where the utility monster starts from a baseline status of massive suffering.

from the philosophical literature dedicated to choices of an optimal length of life under utilitarianism. That literature studied the selection of a length of life to highlight some unattractive implications of utilitarianism. In particular, Cowen (1989) proposed a paradoxical result – the Methuselah Paradox – which is a variant of Parfit’s Repugnant Conclusion, but based on multiplying life-periods rather than lives: “for any possibly ecstatically happy and profound life of, say, 200 years, we can imagine another, much longer life which will welfare-dominate it simply by multiplying many years of epsilon utility” (Cowen 1989, p. 37).⁴

The present article will focus only on the formal relationship between inequalities in longevity – taken as given – and the existence of Nozick’s “utility monsters”. Our analysis will reveal that, although a long-lived person, if taken separately, cannot be a “utility monster”, a large group of long-lived persons can, under some conditions, constitute what can be called a “collective utility monster”, that is, a group of persons that would, under utilitarianism, capture (almost) all resources. Based on these findings, I will then conclude that, in modern societies where a high proportion of the population reaches very old ages, large groups of long-lived persons can, under plausible conditions, be regarded as “collective utility monsters”. After a thorough review, Nozick’s objection against utilitarianism turns out to be most relevant for real-world aging societies. Aging raises a real puzzle for utilitarianism.

I will proceed in four stages. I will first provide a precise account of the “utility monster”. A key question is whether a “utility monster” must have a life that is millions of times better than other lives, as supposed in Parfit (1984). In Section 2, I will argue that such a requirement does not constitute a necessary component of what a “utility monster” is. I will defend a purely marginalist view of “utility monster”, that is, that the primary characteristic of a “utility monster” consists of a higher *marginal* well-being level rather than of a higher *total* well-being level. I will defend this marginalist account of “utility monster” by demonstrating

⁴ For a defence of utilitarianism against the Methuselah Paradox, see Ng (1989).

that it is possible that one person captures almost all resources without having a life that is much better than other lives. Then, in a second stage, I will present three purely marginalist accounts of “utility monsters” (Section 3). While these accounts do some justice to Nozick’s criticism of utilitarianism, they exclude the possibility that a group of persons with a higher sensitivity of well-being captures (almost) all resources under utilitarianism. To allow for the possibility of collective predation, Section 4 will propose three marginalist accounts of a “collective utility monster”, defined as a group of persons who would, under utilitarianism, collectively capture (almost) all resources. In a fourth stage, I will consider a cake-division problem between ordinary persons (enjoying a normal lifespan) and long-lived persons, and I will demonstrate that, although a single long-lived person can hardly be a “utility monster”, a large group of long-lived persons can, under some conditions, constitute a “collective utility monster”. I will finally deduce some implications of this result for real-world aging societies.

2. Utility monsters: a marginalist view

When examining the possibility of existence of “utility monsters”, Parfit (1984) imposes a particular requirement on such “monsters”: it must be the case that their lives exhibit a much higher total well-being than the lives of other persons. Parfit considers that “utility monsters” achieve extremely high well-being levels, that is, that the “Monster’s quality of life must be *millions* of time as high as that of anyone we know.” (Parfit 1984, p. 389).

Parfit’s account of the “utility monster” is not purely marginalist: it does not focus only on the stronger capacity of monsters to transform, at the margin, resources into well-being. According to Parfit, the monster is a being that achieves extremely high levels of well-being. This requirement makes the existence of “utility monsters” hardly plausible. Indeed, it is

difficult to imagine a person whose life would be millions of times better than the lives of other persons. From this difficulty, Parfit concludes that utility monsters cannot exist.

While this conclusion follows logically from Parfit's premises, one can nonetheless question at least one of these premises. Parfit's conclusion about the deep impossibility of "utility monsters" relies on the (implicit) assumption stating that the achievement of extremely high well-being levels (in comparison to other persons) is a necessary component of what a "utility monster" is. That assumption can actually be questioned.

To see why, it is useful to re-examine the context in which Nozick proposed his example of "utility monster". This "utility monster" is introduced to demonstrate that, in the presence of these beings, utilitarianism leads to an unattractive implication: the concentration of (almost) all resources in the "maw" of the utility monster, whereas other persons are sacrificed on the ground of total well-being maximization. It should be stressed that this unattractive result – the concentration of (almost) all resources in the "maw" of the monster – does *not* require that the monster achieves an extremely high level of total well-being, nor a level of well-being much larger than the one of other persons. The only requirement for that result is that the well-being of the monster has a larger reactivity or sensitivity to resources in comparison to other persons. This larger sensitivity may, in some cases, lead the monster to achieve much higher total levels of well-being than other persons, but this does not need to be the case.

It is not true that Nozick's critique of utilitarianism requires that "utility monsters" achieve extremely high levels of well-being. The problem raised by Nozick does not primarily concern the *total* well-being level achieved by the monster, but the *marginal* well-being gains derived by the monster. In different words, the unattractive implications of the monster's case under utilitarianism do not require that a single person has a well-being level that is millions of times higher than the one of other persons. The unattractive distributive implications of the

monster's case under utilitarianism only require that the monster has a utility function that is more sensitive to resources, leading to greater utility gains than other persons.

To illustrate this point, let us consider a simple cake division problem with two persons A and B. Their utility functions $U_A(c_A)$ and $U_B(c_B)$ are represented on Figure 1. Figure 1 shows the distribution of cake (c_A^*, c_B^*) that maximizes the sum of utilities, that is, the distribution that equalizes marginal utilities across persons (what Sen (1980) calls “utilitarian equality”).⁵ In this example, person A receives, under utilitarianism, almost the entire cake, because of a higher marginal utility of the cake for all cake levels. On the contrary, person B receives a much smaller piece of cake (that is, $c_B^* \ll c_A^*$), because of a flatter utility function.




Figure 1. Utilitarian equality in a simple cake division problem.

In that example, person A is a “utility monster”, leading to the sacrifice of person B. But it is not the case that person A’s well-being is many times larger than person B’s well-being. Thus

⁵ On Figure 1, the equalization of marginal utilities at the allocation (c_A^*, c_B^*) is illustrated by the equalization of the two slopes (emphasized by thick lines) of the utility functions $U_A(c_A)$ and $U_B(c_B)$ at the allocation (c_A^*, c_B^*) .

Figure 1 shows an example of a person that does not achieve a total well-being level that is many times higher than that of other persons, but who is nonetheless a “utility monster” in Nozick’s sense. This example illustrates that being a “utility monster” does *not* require the achievement of a life many times better than other lives. Actually, being a “utility monster” only requires that one has well-being that is highly sensitive or reactive to resources.

The primary feature of a “utility monster” is not to achieve an extremely high level of well-being, but, rather, to capture (almost) all resources because of a larger sensitivity of well-being. It is that particular feature – and no other – that explains why utilitarianism allocates (almost) all resources to the monster. That characteristic being decisive, there is a strong case for defining a “utility monster” solely by a larger reactivity of well-being to resources. This view can be called *a purely marginalist view* of the “utility monster”. A purely marginalist account of the “utility monster” focuses only on the sensitivity of the person’s well-being to resources, without any requirement concerning the achieved total well-being level.

Once one accepts this purely marginalist view of “utility monster”, Parfit’s criticism loses some of its strength. The existence of a “utility monster” does not require having persons with well-being achievements much higher than the ones of other persons. True, such persons do not exist in the real world and will probably never exist. But this does not tell us anything about the non-existence of “utility monsters”. Under a purely marginalist view, having a much higher total well-being is not necessary for being a “utility monster”.

3. Three purely marginalist accounts of “utility monster”

As shown in Section 2, it is a particular characteristic of the utility monster – a larger sensitivity of well-being with respect to other persons – that leads to the concentration of

(almost) all resources into the “maw” of the monster, as this was criticized by Nozick (1974). But even if one adopts a purely marginalist account, the definition of a “utility monster” remains incomplete. What does a “larger sensitivity of well-being with respect to other persons” mean? Depending on how these terms are defined, several accounts of the “utility monster” can be distinguished. This section will present three purely marginalist accounts.

Consider a resource in total quantity X to be divided among N persons and one person M , the “utility monster”. Given that the primary characteristic of the monster concerns the sensitivity of well-being to resources, one can hardly define precisely a “utility monster” without specifying individual utility functions. Let us suppose that persons have utility functions $U_i(x_i)$ where x_i denotes the quantity of the resource enjoyed by person i . Utility functions are supposed to be increasing (i.e., the marginal derivative of $U(.)$ is positive, that is, $U_i'(x_i) > 0$) and concave (i.e. $U_i''(x_i) < 0$), that is, there is positive but declining marginal utility.⁶

A first account defines a “utility monster” as a person for whom the capacity, at the margin, to transform resources into well-being is *always* larger than the capacity of any other person to transform resources into well-being, whatever the distribution of the resource is.

Utility monster (account I)

A person M is a utility monster if and only if $U_M'(X) \geq U_j'(0)$ for all $j \neq M$.

Under this account of the “utility monster”, the capacity of the monster to convert resources into well-being is stronger than that of any other person, and so strong that the sacrifice of

⁶ As usual, the derivative $F'(x)$ of a function $F(x)$ is defined as the limit, when Δx tends to 0, of the ratio $[F(x + \Delta x) - F(x)]/\Delta x$. The concept of marginal utility is due to Jevons (1874), who referred to these as the “final degree of utility”. Jevons introduced the law of declining marginal utility.

others takes an extreme form: all resources should, under utilitarianism, go into the “maw” of the monster. The condition mentioned in this first account is sufficient to have all resources given to the monster under utilitarianism. To see why, consider the reduction of the share given to the monster from X to $X - e$, and divide e into N equal pieces e/N to be distributed among the N persons. Given $U_M'(X) \geq U_j'(0)$, we have, by concavity of utility functions, that $U_M'(X - e) > U_j'(e/N)$, implying that transferring the piece e/N from person j to the monster would increase the sum of utilities. Moreover, since $U_M'(X - e + (e/N)) > U_k'(e/N)$, redistributing a second piece e/N from another person k to the monster would also raise the sum of utilities. The same rationale can be reproduced for all pieces e/N . When all pieces e/N except one are transferred to the monster, we have $U_M'(X - e + (N - 1)(e/N)) > U_N'(e/N)$, so that, here again, the sum of utilities is raised by transferring this last piece e/N to the monster.

An important corollary of this is the following: to obtain the outcome where all the resource is given to the monster under utilitarianism, it is not necessary that the marginal utility for the monster when he has all the resource $U_M'(X)$ is larger than the *sum* of marginal utilities for the N other persons $\sum U_j'(0)$. It is sufficient that $U_M'(X) \geq U_j'(0)$ for all j . Even if there are millions of ordinary persons, a condition on the marginal utility of the monster with respect to the marginal utility of each other person suffices so that the monster obtains all the resource.

At this stage, one may wonder which role plays the “0” in the account I of the “utility monster”. Is it necessary, to be described as a “utility monster”, to let no resource available *at all* to other persons? Or, alternatively, is this characterization of a “utility monster” too narrow, by focusing on a non-generic case where the monster captures the entire resource? It can be argued that a person with a large sensitivity of well-being can be defined as a utility monster, even if this sensitivity leads to a utilitarian distribution where other persons receive a strictly positive amount of the resource. Replacing 0 by an extremely small number – such as 0.000000000000001 – does not significantly alter the meaning of being a “utility monster”.

Let us make a thought experiment. Suppose that, when dividing a cake on a utilitarian basis, *almost* the entire cake goes to a given person, except an extremely small amount, which is divided equally among other persons. Let us denote by ε the small piece of cake distributed to each of these persons. In that case, the person who receives almost the entire cake can be regarded as a “utility monster”, even though she does not belong to the category of “utility monster” defined in account I. This intuition is captured by the account II of “utility monster”.

Utility monster (account II)

A person M is a utility monster if and only if there exists a number $\varepsilon \geq 0$ and a constant $C > 1$ with $\varepsilon < X / C(1 + N)$ such that $U_M'(X - N\varepsilon) \geq U_j'(\varepsilon)$ for all $j \neq M$.

In account II of “utility monster”, two parameters are introduced. First, the parameter ε describes a small amount of the resource at which the sensitivity of the monster’s well-being (weakly) exceeds that of other persons. The parameter ε was set to 0 in account I, but ε could be strictly positive, as long as it is sufficiently small. Note that it is difficult to define what “sufficiently small” means. This is the task of the second parameter introduced in account II: the constant C . C specifies a threshold defining how small ε should be in order to have in a situation where the person M is considered to be a “utility monster”. If C were equal to 1, the threshold for a “sufficiently small” ε would coincide with the amount that each person would have under the egalitarian distribution. C being higher than 1, the threshold is a downward deviation from the egalitarian distribution. One can thus call C the “deviation constant”, since the constant C specifies the degree to which the distribution must deviate from the egalitarian distribution to be qualified as a situation where a “utility monster” is at work.

When ε converges to 0, the inequality $\varepsilon < X / C(1 + N)$ is necessarily satisfied and the inequality in marginal utilities in account II coincides exactly with the inequality in marginal utilities stated in account I. Account II of utility monster thus includes, as a special case, account I. But given that Nozick's writings define the "utility monster" without specifying precise inequalities in marginal utilities, there exists no reason to assume that the parameter ε must be fixed to 0. Replacing 0 by an infinitesimally small number still preserves the intuition of what is meant by a "utility monster": a person for whom the welfare gains are so large that the interests of other persons must be sacrificed from a utilitarian point of view.⁷

Whereas account II assumes that the extremely small amount of resource given to each person is equal to a constant ε , there is no reason to focus only on that special case. Consider another thought experiment, where the utilitarian division of the cake gives *almost* the entire cake to a single person, whereas other persons receive extremely small – but unequal – shares. The person receiving almost the entire cake can be regarded as a monster. But she does not fall under the concepts of "utility monster" defined in accounts I and II. Hence, accounts I and II do not fully capture what a "utility monster" is. This justifies introducing account III.

Utility monster (account III)

A person M is a utility monster if and only if there exist numbers $\varepsilon_j \geq 0$ for all $j \neq M$ and a constant $C > 1$ with $(1 / N) \sum_{j \neq M} \varepsilon_j < X / C(1 + N)$ such that:

$$U_M'(X - \varepsilon_1 - \dots - \varepsilon_N) \geq U_j'(\varepsilon_j) \text{ for all } j \neq M.$$

⁷ Account II provides a sufficient condition for having all resources except $N\varepsilon$ given to the monster. It is not necessary that $U_M'(X - N\varepsilon)$ exceeds the *sum* of $U_j'(\varepsilon)$ across the N other persons. The underlying intuition goes as follows. If one reduces the monster's share from $X - N\varepsilon$ to $X - N\varepsilon - e$, and divides e into N pieces distributed to each of the N persons, we have, by concavity of utility functions, that: $U_M'(X - N\varepsilon - e) > U_j'(\varepsilon + (e/N))$, so that transferring the piece e/N to the monster would raise the sum of utilities. The same rationale holds for all pieces e/N . Hence, the condition $U_M'(X - N\varepsilon) \geq U_j'(\varepsilon)$ is sufficient to have predation of almost all resources.

Account I of “utility monster” corresponds to account III when all numbers ε_j are set to 0. Moreover, account II coincides with account III when all ε_j are equal to a constant $\varepsilon \geq 0$. But there is no reason to restrict the occurrence of monsters to these special cases. Nozick’s definition of “utility monster” does not require that the small amounts of resources left to ordinary persons should be equal to zero or equal to each other. The only requirement is that the marginal welfare gains of resource for the monster dominate the welfare gains of other persons, so that only a small proportion of the resource is left out of the “maw” of the monster. This is what account III of the “utility monster” formalizes.⁸

In sum, this section provided three accounts of what a “utility monster” consists of. Under each account, the monster is characterized by a larger sensitivity of well-being to the resource, which makes utilitarianism assign (almost) the entire resource to the monster. The extremely small amount left to ordinary persons is either fixed to 0 (account I), to a small positive amount equal for all (account II), or to a small positive amount that may vary across persons (account III). Given that there is no reason to assume that the resource left out of the “maw” of the monster must be equally divided across all, account III is the most attractive account.

4. Collective utility monsters

Although they capture some aspects of the idea of “utility monster”, accounts I to III suffer from a major limitation. They define the “utility monster” by means of a comparison between

⁸ To see why the condition in account III suffices for the predation of (almost) all resources by a single person, consider a reduction of the monster’s share by an amount e . Whatever the precise way in which this amount e is shared among the N persons, transferring these (possibly unequal) pieces back to the monster would necessary increase the sum of utilities (because of $U_M'(X - \varepsilon_1 - \dots - \varepsilon_N) \geq U_j'(\varepsilon_j)$ and the concavity of utility functions).

a single person with high marginal utility – the monster – and N ordinary persons. This kind of “one against all” comparison does not do entire justice to Nozick’s thought. When developing his criticism of utilitarianism, Nozick considered the possibility of “utility monsters” *in the plural*, and not necessarily a unique person capturing (almost) all resources.

However, accounts I to III are incompatible with the non-uniqueness of the utility monster. To see this, consider a person M with a high marginal utility to the resource. Person M is an obvious candidate for being qualified as a “utility monster”. Then, duplicate person M into person M' . The mere existence of that replica prevents M from being qualified as a “utility monster” under accounts I to III. Indeed, once person M' exists, it is no longer true that M exhibits a higher sensitivity of well-being than *any other person in the population*. Thus, once M' exists, person M can no longer be qualified as a “utility monster” under accounts I to III.

This is an important limitation of accounts I to III of “utility monster”, because there is no reason to exclude the possibility of predation of resources by a *group* of persons with a high sensitivity of well-being. In his objection against utilitarianism, Nozick was concerned with the unattractiveness of the predation of (almost) all resources by some persons. From an ethical perspective, the predation of resources by, let us say, two or three persons, is (nearly) as unattractive as the predation by a single person. Hence, the criticism of the distributive implications of utilitarianism must allow for the predation by groups that are not singletons.

Accounting for collective predation requires to give up the “one against all” framework of Section 3, and to consider instead a setting where the population is partitioned in two groups: on the one hand, a group \mathbf{N} including a number N of ordinary persons; on the other hand, a group \mathbf{O} including a number O of persons with a higher sensitivity of well-being.⁹ Note that,

⁹ The term “partition” implies that no person belongs to *both* groups \mathbf{N} and \mathbf{O} . Moreover, it is assumed that $N \geq O$, in order to distinguish the case of “collective utility monsters” from the case of Parfit’s population Z , where a group would, only because of its extremely large size, capture (almost) all resources under utilitarianism.

since \mathbf{O} includes several persons with high marginal utility of the resource, each person in \mathbf{O} cannot be a “utility monster” under accounts I to III. But even if \mathbf{O} includes no monster, it can be the case that this group captures (almost) all resources under utilitarianism, because of the higher marginal utilities of its members. A group that leaves (almost) nothing to the rest of the population has the ugly appearance of a monster. Hence, when a group captures (almost) all resources under utilitarianism, we call this group a “collective utility monster”.¹⁰ As for “utility monsters”, the definition of “collective utility monster” depends on what one means by “capturing (almost) all resources”. This leads to the following three accounts.

Collective utility monster (account I)

A group \mathbf{O} of O identical persons is a collective utility monster if and only if for all persons m in \mathbf{O} and for each person j in \mathbf{N} , we have: $U_m'(X/O) \geq U_j'(0)$.

Collective utility monster (account II)

A group \mathbf{O} of O identical persons is a collective utility monster if and only if there exists a number $\varepsilon \geq 0$ for all persons j in \mathbf{N} and a constant $C > 1$ with $\varepsilon < X / C(O + N)$ such that, for each person m in \mathbf{O} , and for each person j in \mathbf{N} , we have:

$$U_m'((X - N\varepsilon)/O) \geq U_j'(\varepsilon)$$

¹⁰ A “collective utility monster” differs from a “population composed of utility monsters”. Under the accounts of Section 3, non-unique persons with higher marginal utilities than ordinary persons do not belong to the category of “utility monster”. But their higher marginal utilities can lead to a collective predation of (almost) all resources, explaining the term “collective utility monster”.

Collective utility monster (account III)

A group \mathbf{O} of O identical persons is a collective utility monster if and only if there exist numbers $\varepsilon_j \geq 0$ for each person j in \mathbf{N} and a constant $C > 1$ with $(1/N)\sum_{j \text{ in } \mathbf{N}} \varepsilon_j < X/C(O + N)$ such that, for each person m in \mathbf{O} , and for each person j in \mathbf{N} , we have:

$$U_m'((X - \varepsilon_1 - \dots - \varepsilon_N)/O) \geq U_j'(\varepsilon_j).$$

When the group \mathbf{O} is a singleton (that is, $O = 1$), the three accounts of “collective utility monster” collapse to the three accounts of “utility monsters” of Section 3. Indeed, in that case, the conditions concerning the higher sensitivity of well-being coincide with the ones of the accounts of “utility monsters”. Moreover, when \mathbf{O} is a singleton, the definition of the deviation constant C is also the same as in the accounts of “utility monster”. Thus a collective utility monster with a single member is necessarily composed of a utility monster.

In the general case where the group \mathbf{O} includes more than a single person, the three accounts of “collective utility monster” differ from the corresponding accounts of “utility monster”. A group \mathbf{O} can fall under the concept of a “collective utility monster” even if none of its members is, on her own, a “utility monster”. The reason is that, under each account, the condition to be part of a “collective utility monster” is *weaker* than the condition to be a “utility monster”. The condition to be a “utility monster” concerns the marginal utility of capturing (almost) all resources, whereas the condition to be part of a “collective utility monster” concerns the marginal utility of capturing *a fraction $1/O$ of (almost) all resources*. Given the concavity of utility functions, the condition for membership of a collective utility monster is weaker than the condition to be a utility monster.

5. Utility monsters and the long-lived

In this section, I argue that adopting purely marginalist accounts of “utility monsters” and “collective utility monsters” has crucial implications when examining how utilitarianism distributes resources among persons having unequal lengths of life. I will show that, although long-lived persons do not generally belong to the category of “utility monsters”, large groups of long-lived persons can, in some cases, be regarded as “collective utility monsters”.

For that purpose, this section will proceed in two steps. First, in order to examine the pure effect of longevity differentials on the distribution of resources under utilitarianism, I will consider a simple cake division problem among two groups that differ *only* regarding the length of life of their members, everything else being the same across groups. Then, in a second step, I will re-examine that cake division problem, while allowing for differences not only in longevity across groups, but, also, in the structure of well-being per life-period.

Consider that a resource in total quantity X is to be allocated within a population that is partitioned in two groups: on the one hand, the group \mathbf{N} , which includes N ordinary persons, with a standard length of life whose duration is normalized to unity, and, on the other hand, the group \mathbf{O} , which includes O persons, with lives that have a length that is K times the normal length of life, with $K > 1$. I will call \mathbf{O} the group of “long-lived persons”.¹¹

In order to study the pure effect of longevity differentials, I assume that members of \mathbf{N} and \mathbf{O} differ only regarding their longevity. The lifetime well-being of each person j in \mathbf{N} is given by a utility function $U_j(x_j)$, that is the same for all members of the group: $U_j(x) = U(x)$ for all j in \mathbf{N} . The lifetime well-being of long-lived persons m in \mathbf{O} is the sum, across K periods, of

¹¹ Individual longevitys are supposed to be independent from how resources are allocated. As I stressed in Section 1, relaxing that assumption would lead us to analyse optimal prevention against death, something that goes beyond the topic of this paper.

temporal well-being levels measured by the same utility function $U(\cdot)$, and is thus equal to $K U(x_m/K)$, where x_m is the total amount of resource given to a long-lived person m .

Proposition 1

Consider a population partitioned in groups \mathbf{N} and \mathbf{O} , whose members differ only on their longevity, the members of \mathbf{O} having a longevity advantage K . The group \mathbf{O} is a collective utility monster under accounts II and III if and only if K is higher than $C + (C - 1)(N/O)$, where C is the deviation constant under accounts II and III.

Proof: See the Appendix.

Proposition 1 states a necessary and sufficient condition under which the group of long-lived persons is a “collective utility monster”: the longevity advantage of long-lived persons should be higher than a threshold that is increasing in the deviation constant C and in the number of ordinary persons N , and decreasing in the number of long-lived persons O .¹²

Consider first the case where the group of long-lived persons is a singleton (that is, $O = 1$). In that case, the condition of Proposition 1 is a necessary and sufficient condition for the unique long-lived person to be a “utility monster”. But when the number of ordinary persons N is sufficiently large, the condition in Proposition 1 is not satisfied. Even a person living as long as purely imaginary beings like Lewis’s Methuselah (Lewis 1976), living 969 years, could not satisfy that condition. Thus a single long-lived person cannot be a “utility monster”.

Consider now the general case where the group of long-lived persons is not a singleton ($O > 1$). Notice that the higher the number of long-lived persons O is, the weaker the condition stated in Proposition 1 is. Thus, the condition under which a group of long-lived persons

¹² Proposition 1 deliberately ignores the special case where ordinary persons receive 0, in which case the population \mathbf{O} would belong to the category of “collective utility monster” under both accounts I, II and III.

belongs to the category of “collective utility monster” is *weaker* than the condition under which a single long-lived person is a “utility monster”. Therefore, it is possible that, even if no single long-lived person is a “utility monster” on her own, a sufficiently large group of long-lived persons can, under some circumstances, be a “collective utility monster”. In other words, even if no single centenarian is a “utility monster” on her own, a sufficiently large set of centenarians could fall under the concept of a “collective utility monster”.

The intuition behind this result is that even if, under utilitarianism, a centenarian person obtains more resources than persons with ordinary lifespan, the utilitarian distribution still leaves, in that case, “sufficiently enough” resources to ordinary persons. However, when considering a sufficiently large group of centenarians, things are different: under some conceptions of “sufficiently small” defined by the deviation constant C , it could be the case that utilitarianism leaves “too little” resources to ordinary persons, in which case the group of centenarians then falls under the concept of “collective utility monster”.

What does Proposition 1 tell us about real-world societies? To answer this question, it is necessary to distinguish societies existing before and after the demographic transition (the long-run tendency toward a fall of fertility and mortality to low levels).¹³ Before the demographic transition, long-lived persons were, in relative terms, very few in the society. The number N was much larger than O , making the condition of Proposition 1 not satisfied. But things are different nowadays. Thanks to the improvement of survival conditions, an increasingly large population reaches very old ages. Under these circumstances, the condition of Proposition 1 could be satisfied for some reasonable values of the deviation constant C .

To illustrate this, let us take the example of Australia, where the median age at death is nowadays as large as 84 years for women and 78 years for men (Australian Institute of Health

¹³ On the demographic transition, see Lee (2003).

and Welfare, 2022). If we take this median age at death as a basis to partition the population in two groups **N** and **O** of equal sizes (representing 50 % of the total population), this leads approximately to an (average) longevity advantage of between 25 and 30 years for the long-lived, equal to about 40 % in relative terms (i.e., $K = 1.40$). This implies that the condition of Proposition 1 is satisfied for all levels of the deviation constant C inferior to 1.40.

As a consequence, although long-lived persons are not, taken one-by-one, “utility monsters”, the large population of persons reaching very old ages in modern societies can, under some definitions of what “sufficiently small” means, be regarded as a “collective utility monster”.

Up to now, our analysis assumed that groups **N** and **O** differ only on the longevity of their members, all other aspects being the same. This *ceteris paribus* assumption was made in order to identify the pure effect of longevity differentials on the utilitarian distribution of resources.

Are the results of Proposition 1 robust to relaxing the *ceteris paribus* assumption?

To answer that question, let us now give up the assumption of same temporal utility functions for all persons. Suppose now that each person j in **N** has a utility function $U_j(x_j)$, which differs across members of **N**. Moreover, the temporal utility of long-lived persons is now given by a function $V(\cdot)$, which is increasing and concave, and the same across all life-periods. It is assumed that utility functions $U_j(\cdot)$ do not exhibit a larger sensitivity than the temporal utility function of the long-lived, that is, for any given x and for all j in **N**, we have: $U_j'(x) \leq V'(x)$.¹⁴

Proposition 2

*Consider a population partitioned in two groups **N** and **O**, whose members differ regarding their longevity, the members of **O** having a longevity advantage K , while*

¹⁴ Here again, we ignore the special case where ordinary persons receive 0, in which case the population **O** would belong to the category of “collective utility monster” under both accounts I, II and III.

temporal utility functions are not necessarily the same across groups. If K is higher than $C + (C - 1)(N/O)$, \mathbf{O} is a collective utility monster under accounts II and III.

Proof: See the Appendix.

When one relaxes the assumption of equal temporal utility functions across groups, a sufficiently large longevity advantage K is no longer a necessary condition for a group of long-lived persons to be a “collective utility monster”, unlike what prevailed under Proposition 1. The underlying intuition is the following. When we allow for temporal utility functions to vary across groups, it is possible that ordinary persons have extremely flat temporal utility functions. In that case, the group of long-lived persons would constitute a “collective utility monster”, even if the longevity advantage K of its members is small. This explains why a high longevity advantage is no longer a necessary condition for the group of long-lived persons to be a “collective utility monster” here, unlike in Proposition 1.

Having stressed this difference with respect to Proposition 1, it remains that Proposition 2 has important implications regarding whether or not a group of long-lived persons can constitute a “collective utility monster”. When interpreting Proposition 1, which assumed same temporal utility functions for all persons, I argued that, in contemporary aging societies where a large proportion of the population reaches very old ages, a group of long-lived persons could belong to the category of “collective utility monster”. Proposition 2 states that, even if one assumes different temporal utility functions, a sufficiently large longevity advantage for the long-lived constitutes a sufficient condition, which guarantees that a group of long-lived persons falls under the concept of “collective utility monsters”. As a consequence, if we turn back to the case of modern aging societies, our corollaries of Proposition 1 still prevail under different temporal utility functions. Back to the case of Australia, the prevailing demographic conditions are such that the sufficient condition stated in Proposition 2 is likely to be satisfied,

with the conclusion that the large group of long-lived persons constitutes a “collective utility monster”. This conclusion holds even under distinct temporal utility functions across groups.

6. Conclusions

What do we learn from all this? The contributions of this article are threefold.

First, contrary to Parfit’s (1984) interpretation of the “utility monster”, the primary characteristic of a “utility monster” lies in the larger sensitivity of well-being to resources, rather than in a higher total well-being level. Adopting this purely marginalist view on “utility monsters” weakens the conclusion that these monsters are deeply impossible on the ground that no life can be many times better than other lives, because this property is not a necessary condition for being a “utility monster” under a purely marginalist account.

Second, adapting the accounts of “utility monsters” in such a way as to allow for the collective predation of (almost) all resources by a group of persons having a larger sensitivity of well-being has important implications. A set of persons who, taken one-by-one, would not be counted as “utility monsters”, can nonetheless be, under some conditions, a “collective utility monster”, that is, a group that would capture (almost) all resources under utilitarianism.

Third, assuming that groups of ordinary persons (with standard lifespan) and long-lived persons differ only in terms of longevity of their members, I derived a necessary and sufficient condition under which a group of long-lived persons is a “collective utility monster” (Proposition 1). Although a single long-lived person cannot fall under the concept of “utility monster”, a sufficiently large set of long-lived persons can be a “collective utility monster”. This case was unlikely in societies existing before the demographic transition, but is possible under aging societies characterized by a high proportion of persons reaching very old ages. If

one relaxes the assumption of same temporal utility functions for all persons, the condition of Proposition 1 is no longer necessary, but is sufficient to guarantee that a large group of long-lived persons with a high longevity advantage is a “collective utility monster” (Proposition 2).

In sum, if we turn back to the initial question “Are long-lived persons utility monsters?”, our analysis provides the following answer. When taken separately, each long-lived person does not fall under the concept of “utility monster”. However, a large group of long-lived persons with a sufficiently high longevity advantage can constitute a “collective utility monster”. If true, Nozick’s objection to utilitarianism would be most relevant for modern aging societies. Aging raises a real puzzle for utilitarianism.¹⁵

7. References

- Australian Institute of Health and Welfare (2022). Deaths in Australia, AIHW, January 2022.
- Briggs, Rachael, & Daniel Nolan (2015). Utility Monsters for the Fission Age. *Pacific Philosophical Quarterly*, 96, 392-407.
- Chappel, Richard Yetter (2021). Negative Utility Monsters. *Utilitas* 33 (4), 417-421.
- Cowen, T. (1989). Normative Population Theory. *Social Choice and Welfare*, 6, 33-43.
- Jevons, W.S. (1874). *The Theory of Political Economy*. New edition, Pellican, London.
- Lee, Ronald (2003). The demographic transition: three centuries of fundamental change. *Journal of Economic Perspectives*, 17, 167-190.

¹⁵ One could argue that in real-world economies, the total quantity of resource is not fixed, but can vary with the demographic structure. Assuming constant resources is thus a simplification. But it should be stressed that reforms aimed at adapting societies to the aging process (such as the postponement of the retirement age) face serious oppositions. The distributive challenge raised by aging is a real one, not a theoretical artefact.

Lewis, David (1976). Survival and Identity. In: A.O. Rorty (ed.), *The Identities of Persons*, University of California Press, Berkeley.

Ng, Y. K. (1989). Hurka's Gamble and Methuselah's Paradox. A Response to Cowen on Normative Population Theory. *Social Choice and Welfare*, 6, 45-49.

Nozick, Robert (1974). *Anarchy, State and Utopia*. Blackwell, Oxford.

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

Sen, Amartya (1980). Equality of What? *Tanner Lectures on Human Values*.

8. Appendix

Proof of Proposition 1

The population being partitioned into a group **N** including N ordinary persons and a group **O** including O long-lived persons, the utilitarian (interior) optimum is a list $\{x_a^*, x_b^*, \dots, x_N^*, x_m^*, x_n^*, \dots, x_O^*\}$. At the utilitarian optimum, there is equality of all marginal utilities associated to the last unit of resource given to all persons:

$$U_a'(x_a^*) = U_b'(x_b^*) = \dots = U_N'(x_N^*) = U_m'(x_m^*) = U_n'(x_n^*) = \dots = U_O'(x_O^*)$$

Under the assumption of same temporal utility functions for all persons, we have: $U'(x_a^*) = U'(x_b^*) = \dots = U'(x_N^*)$, which implies: $x_a^* = x_b^* = \dots = x_N^*$. In the rest of the proof, we will take person N as representative of all ordinary persons.

For long-lived persons, $U_m(x_m) = \dots = U_O(x_O) = KU(x_m/K)$, so that, taking person O as a representative of all long-lived persons, we have: $U_O'(x_O^*) = KU'(x_O^*/K)(1/K) = U'(x_O^*/K)$. Given that, at the utilitarian optimum, we have $U'(x_N^*) = U_O'(x_O^*) = U'(x_O^*/K)$, it follows

that $x_O^* = Kx_N^*$. Note also that $X = Nx_N^* + Ox_O^*$. Substituting for $x_O^* = Kx_N^*$, we obtain that: $Nx_N^* = X - O(Kx_N^*)$, which implies $x_N^* = X / (OK + N)$.

Sufficiency. Let us prove that, if $K > C + (C - 1)(N/O)$, the group **O** is a collective utility monster. Under account III, **O** is a collective utility monster if and only if there exist numbers $\varepsilon_j \geq 0$ for each person j in **N** and a constant $C > 1$ with $(1/N)\sum_{j \text{ in } N} \varepsilon_j < X / C(O + N)$ such that, for each m in **O**, and for each j in **N**, we have: $U_m'((X - \varepsilon_1 - \dots - \varepsilon_N)/O) \geq U_j'(\varepsilon_j)$.

To show that this condition is satisfied when $K > C + (C - 1)(N/O)$, let us set: $\varepsilon_a = x_a^*$, $\varepsilon_b = x_b^*$, ..., $\varepsilon_N = x_N^*$. Substituting for all ε_j in equations for marginal utilities, we obtain:

$$U_a'(\varepsilon_a) = \dots = U_N'(\varepsilon_N) = U_m'((X - \varepsilon_a - \dots - \varepsilon_N)/O) = \dots = U_O'((X - \varepsilon_a - \dots - \varepsilon_N)/O)$$

These conditions characterize a “collective utility monster” if the condition $(1/N)\sum_{j \text{ in } N} \varepsilon_j < X / C(O + N)$ is satisfied. Let us prove this. We have, at the utilitarian optimum:

$$(1/N)\sum_{j \text{ in } N} \varepsilon_j = X / (OK + N)$$

When $K > C + (C - 1)(N/O)$, we have $X / (OK + N) < X / C(O + N)$, so that:

$$(1/N)\sum_{j \text{ in } N} \varepsilon_j < X / C(O + N)$$

as required for **O** to be a collective monster under account III. Finally, note that the ε_j terms also satisfy the conditions guaranteeing that **O** is a collective utility monster under account II.

Necessity. Let us show that, if the group **O** is a collective utility monster, $K > C + (C - 1)(N/O)$ holds. To prove this by *reduction ad absurdum*, assume $K \leq C + (C - 1)(N/O)$.

When **O** is a collective utility monster, it has to be the case, at the utilitarian optimum, that the x_j^* satisfying utilitarian equality (equalization of marginal utilities at the optimum) are equal to ε_j that are sufficiently small, that is, such that: $(1/N)\sum_{j \text{ in } N} \varepsilon_j < X / C(O + N)$.

At the utilitarian optimum, we have, as shown above that for all persons j in **N**, $x_j^* = x_N^* = \varepsilon_N$, and for all persons m in **O**, that $x_m^* = Kx_N^* = K\varepsilon_N$. Therefore: $\varepsilon_N = X / (OK + N)$.

When $K \leq C + (C - 1)(N/O)$, we have $X / (OK + N) \geq X / C(O + N)$, so that:

$$\varepsilon_N = X / (OK + N) \geq X / C(O + N)$$

This contradicts the condition: $(1/N) \sum_{j \in N} \varepsilon_j < X / C(O + N)$ to be a collective utility monster.

A contradiction is reached. Thus $K \leq C + (C - 1)(N/O)$ does not hold.

Proof of Proposition 2

At the utilitarian (interior) optimum $\{x_a^*, x_b^*, \dots, x_N^*, x_m^*, x_n^*, \dots, x_O^*\}$, we have:

$$U_a'(x_a^*) = U_b'(x_b^*) = \dots = U_N'(x_N^*) = U_m'(x_m^*) = U_n'(x_n^*) = \dots = U_O'(x_O^*)$$

Since all persons in population **O** have the same utility function, we can base our analysis on a representative long-lived person m with total utility $U_m(x_m) = K V_m(x_m / K)$. By assumption, we have that: $U_j'(x) \leq V_m'(x)$ for all x and all persons j in **N**. Hence, at the utilitarian optimum, we have also that $x_j^* < x_m^* / K$, which implies that $\sum_{j \in N} x_j^* < N x_m^* / K$. We thus have: $X - O x_m^* = \sum_{j \in N} x_j^*$. Hence $x_m^* = (X - \sum_{j \in N} x_j^*) / O$.

Hence, substituting for x_m^* , we have that:

$$\sum_{j \in N} x_j^* < N x_m^* / K = (N/KO) (X - \sum_{j \in N} x_j^*).$$

It follows from this that: $\sum_{j \in N} x_j^* < X (N / (KO + N))$.

In order to show that **O** is a collective utility monster when $K > C + (C - 1)(N/O)$, let us set $x_a^* = \varepsilon_a, \dots, x_N^* = \varepsilon_N$, and substitute for x_m^* in the above equations of utilitarian equality:

$$U_a'(\varepsilon_a) = \dots = U_N'(\varepsilon_N) = U_m'((X - \varepsilon_a - \dots - \varepsilon_N)/O) = \dots = U_O'((X - \varepsilon_a - \dots - \varepsilon_N)/O)$$

These equalities characterize a collective utility monster under account III, provided $\varepsilon_a, \dots, \varepsilon_N$ are small numbers, that is, provided: $(1/N) \sum_{j \in N} \varepsilon_j^* < X (1/C(O + N))$.

Let us check this. At the utilitarian optimum, we have $\sum_{j \in N} \varepsilon_j^* < X(N/(KO + N))$. Moreover, when $K > C + (C - 1)(N/O)$, we have also $X/(KO + N) < X/C(O + N)$. Hence:

$$(1/N) \sum_{j \in N} \varepsilon_j^* < X/C(O + N)$$

as required to have **O** being a collective utility monster under account III. Finally note that the terms ε_j satisfy also all conditions for **O** to be a collective utility monster under account II.