

The social disvalue of premature deaths

Hilary Greaves

January 8, 2014

Biographical information: Hilary Greaves is Fellow and Tutor in Philosophy at Somerville College, Oxford. Her research interests include population ethics, the consequentialist/nonconsequentialist debate, various issues lying at the intersection of ethics and economics, and formal epistemology.

Abstract

Much public policy analysis requires us to place a monetary value on the badness of a premature human death. Currently dominant approaches to determining this ‘value of a life’ focus exclusively on the ‘self-regarding’ value of life — that is, the value of a person’s life to the person whose death is in question — and altogether ignore effects on other people. This procedure would be justified if, as seems *intuitively* plausible, other-regarding effects were negligible in comparison with self-regarding ones. I argue that in the light of the issue of overpopulation, that intuitively plausible condition is at best highly questionable. Unless the world is in fact *underpopulated*, the social disvalue of a premature death is likely to be significantly lower than the current estimates.

Keywords: value of a life, death, overpopulation, underpopulation

There is more to be said for rough estimates of the precise concept than precise estimates of economically irrelevant concepts. (Mishan, 1971)

Many public decisions confront us with the daunting task of placing a monetary value on the badness of a premature human death. A given road safety project would reduce road deaths by a well-estimated number and would cost a known amount; we must decide whether it is cost-effective or not. Increasing the overall budget of the NHS would increase the number of lives that service could afford to prolong, but at some point we have to stop and save resources for other public policy goals. More globally, cutting greenhouse gas emissions, and/or investing in capture and storage projects, would reduce the number of lives cut short by climate-change-induced disasters by figures that can at least be guessed at; the value placed on the badness of premature deaths can be one of the factors significantly affecting recommendations for how much mitigative action is cost-effective.

Many people find the task distasteful. There are two ways of attempting to give intellectual content to this distaste. The first is to assert that human life has a value

that is *lexically prior* to that of mere worldly goods — that is, that given any such choice, the option that saves more human lives is always better, no matter what the cost. The second is to assert that human life is *incommensurable* with worldly goods — that is, that given a choice between (say) reducing the number of road deaths by three and having an extra £4m to spend on other public services, neither option is better or worse than the other, and it is also not the case that the two options are equally good; that they literally cannot be placed in any common ranking. But in the context of public decision making, the ‘lexical priority’ claim leads to absurd results, while the ‘incommensurability’ claim is mere refusal to face up to the issue. Rational public decision-making requires that we settle on a number. What, though, should the number be, and which factors are relevant to determining it?

Several techniques exist for estimating the monetary value of a life, but none is satisfactory. Early approaches were *other-regarding*: that is, they focussed on the importance of the lost life *to those who remain living*. These approaches were roundly criticised in the 1970s, and have since largely been displaced by *self-regarding* approaches, i.e. approaches whose primary focus is the value of the lost life *to the person whose death is in question*. John Broome apparently agrees that this move from other- to self-regarding estimates is a step in the right direction, but has been particularly vocal in criticising the ‘willingness-to-pay’-based approach that is dominant within the latter tradition. In recent work, Broome sketches a (self-regarding) alternative, based on a variant of the ‘quality-adjusted life year’ concept well known to health economists, and intended to keep the merits of the WTP approach while avoiding its drawbacks.

The central point of the present paper is that neither an exclusively other-regarding nor an exclusively self-regarding approach has much hope of being even approximately accurate. Here as elsewhere, the value of an event is determined by its effects on *everybody*, not only by its effects on the person who is intuitively the subject of that event. Here as elsewhere, the sums of very small effects on very large numbers of other people can be comparable in magnitude to the sums of large effects on the small number of people most obviously affected by the event in question. These observations would be mere pedantry if in fact either self- or other-regarding effects dominated the calculation. But I do not think we can be at all confident that that is the case (in either direction). In particular, when we try to square our thinking about the badness of premature deaths with our thinking about overpopulation, we must at least take very seriously the possibility that other-regarding and self-regarding effects are indeed comparable in magnitude. It follows that neither can safely be ignored. An unfortunate consequence is that the correct calculation is highly complex, and that therefore we are far more uncertain about the appropriate monetary figure to place on the disvalue of a premature death than any of the current literature acknowledges. This is a formidable problem for practical decision-making (in particular, in the context of climate change). But we must not pretend that the keys are under the lamppost.

The structure of the paper is as follows. Section 1 briefly reviews the extant approaches to determining the monetary disvalue of a premature death. Section 2 is an exposition of Broome’s critiques of the willingness-to-pay approach, and of his modified-QALY suggestion. Section 3 attempts a relatively complete list of the factors that contribute to the badness of a premature death, and comments on which factors are taken into account by which extant approaches. Section 4 contains the main point of the

paper: I argue that plausible thinking about overpopulation entails that neither other- nor self-regarding aspects of the disvalue-of-death calculation are negligible. Section 5 summarises, and briefly considers how we might proceed in practice in the light of the issues raised in this paper.

1 Extant approaches to valuing life

The primary other-regarding approaches are based on some estimate of a person's output. The *gross output* approach takes the disvalue of a premature death to consist in the loss of the contributions that the deceased would have made to society had he lived out the remainder of his natural life. Average discounted lifetime earnings are used as a crude estimate of this quantity, the rationale being that this figure is "what the society is willing to pay the individual (and thus the person's 'marginal product')" (Cline, 1992, p.117). A (dramatically different) variant on the gross output approach is the *net output* approach. Here, the monetary disvalue of a premature death is taken to be given by average lifetime earnings *less average lifetime consumption*.

The gross output approach dominated until the 1970s.¹ A curious feature of any other-regarding approach, however, is that it altogether ignores the *victim's* interest in living longer. In part due to misgivings about this omission, and in part due to concerns about the shakiness of the above methods' theoretical foundations (Mishan, 1971), they have since largely been displaced by the *willingness to pay* and (occasionally) *hedonic pricing* approaches.

Hedonic pricing techniques attempt to infer the value people place on a non-market good from their revealed willingness to pay for a suitably related market good. For example, householders' implicit valuations of clean air can be derived from the difference in house prices between areas that have different levels of air pollution but are otherwise similar. In the present case, the idea is to derive implicit value-of-life figures from wage differentials: the extra amount a worker must be paid to render him willing to take on a riskier job tells us something about the value he places on his life.

Economists generally regard revealed-preference techniques as superior in principle to stated-preference techniques, on the grounds that what people *do* is a more reliable guide to their true preferences than what they *say* they would do in as yet unencountered situations. However, revealed-preference techniques often face significant methodological problems. In the example just given, for instance, it has been objected that labour market imperfections, such as unionisation, prevent wage differentials from revealing workers' revaluations of the risk of death. This brings us to the currently dominant stated-preference approach to valuing life, by willingness to pay surveys.

¹It appears still to dominate in the developing world. A relatively recent report of the Asian Development Bank (2003) remarks, intriguingly, that

"Since the late 1980s willingness to pay has increasingly been applied for accident costing in industrialized countries. However, the gross output method has been the most commonly used method in most countries over the past few decades. This method is recommended for developing countries as its primary objective, increasing a country's wealth, is thought more appropriate to their needs" (*ibid.*, p.2),

but does not elaborate further.

In such a survey, respondents report the maximum amount they would be willing to pay (or the minimum amount they would be willing to accept) to have a given project go ahead. The total willingness-to-pay (WTP), summed over all affected parties, is then usually taken to be the monetary value of the project in question. This technique is widely applied to, for example, the valuation of environmental amenities ('how much would you be willing to pay to avoid having the size of your local nature reserve decreased by 30%'?).

There is an obvious problem in applying this technique to the case of human life, i.e. the case in which the 'project' is causing or averting the death of the respondent: in that case the relevant WTA figure (at any rate) may well be infinite. The theorist is then threatened with the conclusion that the 'value of a life' for the purposes of public-policy analysis is also infinite. But accepting an infinite 'value of a life' would amount to granting life lexical priority over ordinary goods; we have already noted that such lexical-priority claims are implausible.

The standard response to this is to drop talk of the 'value of a life' altogether, replacing it with talk of the value of decreasing *risk* to life, since WTP/WTA figures for the latter are finite and well-defined. Dividing the result by the change in risk involved, one obtains the so-called 'value of a statistical life' (VOSL). One is then supposed to use the resulting VOSL figures in place of a monetary value for an actual life, bearing in mind that the latter concept is strictly meaningless.

2 Broome on the disvalue of a premature death

2.1 Against WTP

Broome has trenchantly criticised the application of willingness-to-pay methodology to the issue of the value of life. There are essentially two independent issues. The first concerns the consequences of applying the WTP approach in the particular way it is in fact (but need not be) normally practised; the second concerns an unavoidable problem with applying the WTP methodology to the issue of the respondent's death.

The first point arises if the WTP-theorist declines to weight WTP/WTA data by respondents marginal utilities of money. This leads to the 'VOSL' approach outlined above. Against this approach, Broome (i) objects to the absence of marginal-utility weightings in cost-benefit analysis in general (e.g. (2004, p.262)), (ii) points out that the above approach has the implausible conclusion that whether or not *we happen to know who* will be killed by some project makes a vast difference to the projects balance of benefits over costs, holding fixed both the facts about who will be killed and our knowledge of the numbers of people killed (1978); (iii) emphasises the radical implausibility of holding that it is only *risk* to life, and not actual life, that has value (2012b).

The need to resort to talk of the 'value of a statistical life' is avoided by the alternative approach, which thinks of the 'costs' and 'benefits' of cost-benefit analysis in terms of utility, rather than money. On that approach, WTP figures are weighted by marginal utilities of money before being summed in the analysis. This avoids the unwanted conclusion that the value of a life is infinite, since the marginal utility corresponding to an

infinite WTA is zero (if the WTA is infinite, that is precisely because money is useless to the (selfish) agent in question when she is dead). WTP for reductions in risk of death can then be used to calculate the *utility* of a life, *on the assumption that* in determining their WTP for such changes, individuals conform to expected utility theory. This procedure is sound in principle, but in practice the latter assumption is implausible: as is well known, individuals are notoriously irrational in their dealings with risk.

It is in response to this sort of problem that Broome suggests abandoning the WTP approach altogether, in favour of an alternative that seeks to elicit valuations from individuals within a framework that *imposes* certain rationality restrictions. A sketch of Broome's alternative requires a brief digression into the health economists' notion of a quality-adjusted life year, or QALY.

2.2 QALYs and WALYs

NHS cost-effectiveness decisions are determined by a metric known as the 'QALY', or 'quality-adjusted life year'. A given state of health is indexed by a 'health utility', supposed to indicate the health-related quality of life enjoyed by a person in that state of health. A health utility of 1 signifies a state of perfect health, 0 death; negative numbers signify states in which the health-related quality of life is worse than being dead. At any stage of a person's life, he has a certain number of remaining QALYs: the time-integral of his quality of life, over the duration of his remaining life (no matter that this number cannot be known with anything approaching certainty). Healthcare interventions - for example, cancer treatments and eye surgery - are assessed for the number of additional QALYs they are expected to give the patient. The scale is cardinalised so that (surveys and focus group discussions suggest that) a typical individual patient would be indifferent between any two futures that held the same number of QALYs for him. For example, if a state of blindness has a health utility of 0.5, a typical individual would be indifferent between twenty years of life in that state of blindness (followed by death) and ten years of life in a state of perfect health (followed by death).² Appropriate cardinalisation is of course crucial: arbitrary monotonic transformations of the scale would radically affect the cost-effectiveness judgments delivered by the QALY approach.

With the cardinalisation of the QALY scale thus fixed, the idea is that the overall goal of the NHS is to maximise the number of additional QALYs given to the population as a whole. Hence, given a fixed overall budget and subject to additional factors such as the quality of the available evidence for the "cost per QALY" estimate, the NHS will *tend* to provide all and only treatments whose cost per QALY lies below a certain threshold. The actual value of this threshold is around £20–30k (National Institute for Health and Clinical Excellence, 2008, section 4.2).

The key for our purposes is that in principle, the QALY concept can also be used to place a monetary value on premature death. The point is that the associated "health

²Other methods of cardinalisation are also possible: for example, respondents may be asked instead ('standard gamble' questions) what *risk* of immediate death they would be willing to tolerate to have the condition cured, rather than ('time tradeoff' questions) how *many years* of life they would be willing to sacrifice in order to have the condition cured. For a critical survey of cardinalisation methods, see (Nord, 1992).

utilities” place changes between the state of being dead and various states of health on a common scale, appropriately cardinalised, with more trivial changes that are easier to value. For example, a lower bound on the value of curing a headache can be obtained from observed willingness to pay for headache-curing medications. Valuations of these more trivial changes can then be scaled up to yield the value of prolonging a (given) life.³

Broome regards this approach to determining the ‘value of a life’ as a great improvement over the WTP approach: whereas the latter takes individuals’ WTP for changes in risk at face value despite the known irrationalities in the responses of ordinary people to decisions involving risk, the former imposes far more structure from the outside (the latter approach ‘applies no theory’ (2004, p.263); see also Broome’s (1994, 1996)). For instance, it correctly imposes the assumption that, other things being equal, it is better to avert someone’s premature death at age 40 than it is at age 80, despite the fact that an 80-year-old might be willing to pay just as much for increased safety as a 40-year-old. Also, the structure of its questioning techniques forces people to accept the reality of the tradeoffs between quantity and quality of life. Broome therefore regards it as a plausible way of treading the appropriate path between the dual spectres of (i) altogether ignoring the actual preferences of real people and (ii) infecting public policy analysis with the irrationalities and ill-informedness to which survey respondents may be subject.

In his (2004), Broome (however) criticises the QALY-based approach to valuing life approach for taking only health into account (see also (Broome, n.d.)). A better approach, he suggests, would adopt the same general structure, but would assign utilities to states of life more widely construed, not to states of health alone. The result would be a utility scale enabling us to quantify the badness of premature deaths in terms of the known badnesses of *any* of the things that affect well-being (so we could use, e.g., the utility saltuses corresponding to salary increases to find the money equivalents), rather than only the goodness/badness of changes in health state (and corresponding WTP for medicines). We might call this the WALY approach (‘well-being-adjusted life years’).

In fact, for the purposes of estimating the self-regarding disvalue of a premature death⁴, I do not think that there is anything *fundamentally* wrong with the unmodified QALY approach. To place a monetary value on the badness of premature deaths, we require an appropriately cardinalised scale that enables us to weigh death-life transitions against *some* other transitions that are easier to value. But in principle, we only need *one* easier-to-value transition; and a health-related one (curing a headache) is as good as any other. In practice, of course, robustness is served by applying the QALY-style survey methodology to a larger number of transitions. For the purposes of this paper, in any case, the differences between the QALY and WALY approaches are minor.

³Perversely, figures for the ‘value of a statistical life’ derived from the WTP approach have also been used to estimate the monetary value of a QALY, and hence of lesser health benefits (e.g. (Mason, Jones-Lee, & Donaldson, 2009)).

⁴My purpose in this paper is narrower than Broome’s purpose in his (2004); he is discussing a comprehensive theory of well-being, whereas I am only discussing the monetary disvalue of a premature death. I agree that WALYs rather than QALYs are required for the former purpose. I am not sure whether or not Broome thinks they are also required for the latter purpose.

3 Social disvalue

Let us allow that the QALY/WALY approach constitutes a reasonable methodology for assessing the disvalue of a premature death *to the person who dies* — the ‘self-regarding’ disvalue of that death. It does not follow that a figure obtained from QALY/WALY surveys is a reasonable figure to use as the ‘monetary value of a life’ in our public policy decisions. The basic point is familiar. In assessing the *social* value of any event, we need to consider its effects, not only on those most directly and obviously affected, but on *everyone*. The parties to a transaction are not always those most affected by it, and even if they are, *its small effects on large numbers of people may add up to a larger overall contribution to social value than its large effects on the most directly affected parties*.

Granted that they *may*, do they in the present case? To begin addressing this question, we must identify the factors that contribute to the total social disvalue of a premature death. In doing so, it will prove useful to separate *deprivation factors* from *transition factors*. By (stipulative) definition, deprivation factors are those aspects of disvalue that are the result purely of the *absence* of the later part of a person’s life; transition factors are those resulting from the *transition* from the existence to the nonexistence of the person in question. The difference is that deprivation factors but not transition factors would still be present if the person in question had never been born at all.

In the following list of the components (positive and negative) of the value of a premature death, deprivation factors are labelled ‘D’, transition factors ‘T’:

- D1** The loss of the benefit that the dead person would himself have derived from his continued life.
- X1** The pain (physical and/or psychological) that the dying person experiences during the actual process of dying, as compared to the scenario in which the person ceases to exist at the same time but undergoes no pain beforehand.⁵
- T1** Increased fear of death and/or fear of certain activities, and the associated changes of lifestyle, in those who remain living. For example, if Joe dies paragliding, this may cause other people who would otherwise have paraglided to refrain from doing so, thus missing out on enjoyment and life-enriching adventure; if Joe is murdered on the street, this may cause others generally to hide indoors more.
- D2** The positive effect of the death due to the freeing-up of the resources (food, water, physical space etc.) that the deceased would have consumed in his remaining time, had he lived a life of normal duration.
- T2** The negative emotional impact on family associated specifically with *bereavement* (*not* with the foregone pleasure in the deceased’s company, for which latter see (T4)).

⁵This is not really an aspect of the badness *of death* at all: *it could have occurred without the death, and would have been equally bad if succeeded by survival rather than death of the person in question.*

- T3** The negative financial and logistical impact associated with the fact that the deceased's family had organised their affairs on the assumption that the deceased would live significantly longer. This includes, for example, the effects on well-being of the family having adjusted to a lifestyle that is no longer sustainable in the absence of a deceased 'breadwinner'.⁶
- T4** The negative impact on partners due to relationship loss. The key idea here is that, on average, a person in a voluntary long-term partnership is better off (both financially and emotionally) than that same person would have been if single. The effect is in practice likely to be almost entirely transitional: while bereavement is likely to be followed by at least some period of single status, if the deceased had never existed in the first place, most partners would have formed a stable relationship with some other partner in the first place, at roughly the same time as the formation of their actual relationship.
- D3** The economic contribution that the deceased would have made to society via the remainder of his working life.
- T5** Economic losses associated with the *disruption* involved in the employer's having to replace the deceased with a new employee (for example, duplicate recruitment and training costs).

This list is not intended to be exhaustive.⁷ But having a *relatively* complete list on the table throws the incompleteness of each of the extant approaches into sharp relief. The 'gross output' approach (surveyed briefly in section 1) appears to be an attempt to estimate factor (D3). The 'net output' approach appears to be an attempt to estimate the algebraic sum of (D3) and (D2). Which factors the WTP approach includes obviously depends on which actors respondents are taking account when determining their responses, but it seems likely that the majority of respondents will be primarily influenced by factors (D1), (T2), (T3) and (D3). (Meanwhile, if the 'value of a statistical life' rhetoric is taken at face value, it seems to suggest that only something like our factor (X1) is taken into account: the disutility of *fear of death*.) Similar remarks apply to the hedonic pricing approach. The QALY and WALY approaches appear to be attempts to measure factor (D1) alone.⁸

It would, of course, be defensible to use an estimate of (D1) as the social disvalue of a premature death *if* other factors were negligible in comparison to (D1). Charity suggests attributing this hypothesis to advocates of self-regarding approaches. I agree that it is a *prima facie* highly plausible hypothesis. However, I do not think we can be

⁶The family is of course also adversely affected by the loss of the deceased's income, quite aside from issues of adjustment. However, this aspect of disvalue is accounted for in factor (D3) below.

⁷In particular, it omits the knock-on effects due to the fact that averting the premature death of a young person often has the consequence that additional people — the children, grandchildren etc. of the 'saved' person — are brought into existence. These knock-on effects can be net positive or negative — that depends on whether the world is currently under- or over-populated — but in either case, they are potentially very large in magnitude, and may well actually dominate the calculation.

⁸A relatively complete approach, driven by the same concern to calculate the full *social* value as that driving the present paper, is sketched in (Arthur, 1981). Thanks to John Broome for alerting me to this article.

at all confident that it is correct. The purpose of the next section is to make the case for this scepticism.

4 Death and overpopulation

Many believe that the world is currently overpopulated. The thought is that the planet's capacity to provide food, water, energy, materials and simple living space, in addition to current artificial infrastructure, is being stretched far beyond optimum usage by the sheer number of people currently living on the planet, with detrimental consequences both for the long-term capacity of the planet to sustain life and for the quality of life of those who do live in the meantime. This is so to such an extent (the thought continues) that it would be better, all things considered, to reduce world population. Thus campaigners urge, for example, increased spending on family planning programmes, and/or voluntary reductions in family size.⁹

However, there is of course another way to reduce world population: via the premature deaths of some of those who already exist. But even the aforementioned campaigners do not advocate deliberate reductions in public safety on these grounds. And even those who do not think it is *true* that it would be good to reduce the birth rate tend to agree that it is *consistent* to hold that this is good while simultaneously holding that additional premature deaths would be bad. This raises the question of how to render consistent the judgment that additional births would be (on balance) bad with the judgment that additional premature deaths would *also* be (on balance) bad: the background theory cannot be as simple as "any transition to fewer people would be an improvement", nor "any transition to more people would be an improvement".

This question forces us to tangle with population axiology. This is in a sense unfortunate, as there is a notorious lack of consensus in population axiology. It is, however, inevitable: if we want to compare the valuation of premature death with the ethics of variable population size, we need to tangle with the ethics of variable population size.¹⁰ We cannot proceed further without making some controversial theoretical assumptions, at least for the sake of argument. I will first make the ones that I think are true, and later consider how the arguments are affected if we assume instead the axiology defended in Broome's (2004).

As above, let us call the quantity captured by factor (D1) the *self-regarding* value of a life: it is the value of the life *to the person whose life it is*. Define the *social value* of a given person's life (in a given context) to be the difference in value *to society as a whole* (or: value *simpliciter*) between the state of affairs in which that life is lived, and a state of affairs in which the person in question never exists but the well-being levels

⁹Academic discussions of the case for these recommendations is thin on the ground. The campaigners' views are propagated e.g. via the websites of World Overpopulation Awareness <www.overpopulation.org>, and Population Matters <www.populationmatters.org>.

¹⁰In fact the two topics are related independently of such explicit comparisons, since, as noted above (fn. 7), premature deaths of people who are not yet past reproductive age affect the future population. This point has been noted many times, not least by Broome (1985). My point in the present section concerns instead the *additional* connection that remains *even if* those who die prematurely would not otherwise have gone on to have (additional) children.

of all other people are (somehow) unchanged. Our controversial theoretical assumption is that the social value of a given life is equal to the self-regarding value of that life.

Suppose further, merely for illustrative purposes, that the world is currently at optimum population, in the following sense: on average, the difference in value between the actual state of affairs and a state of affairs in which a randomly selected actual person does not exist and others' lives go *as they would have if the person in question had not existed* is zero. It follows that the (positive) self-regarding value of an average life is exactly balanced by the net negative effect of the average person's existence on other people. But the latter quantity is the sum of other-regarding deprivation factors. Thus, under our current assumption, if the world is currently anywhere near optimum population then other-regarding effects are certainly not negligible. It also follows that, under the same conditions, the social disvalue of a premature *death* is the (negative of the) sum of *transition factors alone*; note that none of the extant approaches accords these effects any significant role. If the world is currently *overpopulated*, of course, then our conclusion about the significance of other-regarding effects, including in particular those relating to the creation of additional people, is only exacerbated. (We will take up the possibility of underpopulation shortly.)

Our assumption that the self-regarding value of a life is equal to its social value may sound innocuous: if social value is merely an aggregation of values to individual people, and if no other people are affected, what else can the value of a change that affects one person reasonably be? I think this is in fact correct. It is, however, a controversial piece of population axiology.¹¹ The axiology that satisfies our assumption is *total utilitarianism*, according to which the value of a state of affairs is given by

$$V_{TU}(X) = \sum_{i \in P(X)} u_i(X). \quad (1)$$

(Here, $u_i(X)$ is the lifetime well-being¹² of person i in state of affairs X , and $P(X)$ is the set of people who exist in X .)

Our assumption is, however, denied by every other axiology. This is not the place for a comprehensive survey of the alternatives, but let us consider in particular the 'critical level utilitarianism' that is advocated by Blackorby, Bossert and Donaldson (Blackorby & Donaldson, 1984; Blackorby, Bossert, & Donaldson, 1995; Charles Blackorby & Donaldson, 2005), and to which Broome gives cautious approval in (Broome, 2004). According to critical level utilitarianism, the value of a state of affairs is given

¹¹It is also *somewhat* controversial even aside from issues of variable population. For example, a value function that applies a concave 'prioritarian' transformation to *individual's von Neumann-Morgenstern utilities* notoriously incorporates a divergence between what is good *for individuals* and what is good *simpliciter*. I have argued elsewhere (? ?) that this is not the right way of formalising the prioritarian intuition. Be that as it may, the issues of population axiology on the one hand, and of which is the correct value function even for cases of fixed population, are essentially orthogonal. My focus here is on the former, and I presuppose a utilitarian solution to the latter for simplicity of exposition.

¹²I use the term 'lifetime well-being' in the same sense as does Broome (2004). Lifetime well-being is a measure of how well a person's life, as a whole, goes for him. It is taken to lie on a ratio scale.

There is some subtlety over how to define the zero point. Broome defines it as a life lived throughout in such a state such that, at any moment, death at that moment would be exactly as good as continued life for the person in question. I have no objection to this definition.

by

$$V_{CLU}(X) = \sum_{i \in P(X)} (u_i(X) - \alpha), \quad (2)$$

where α is some constant (generally assumed to be non-negative). According to critical level utilitarianism, therefore, the social value of an additional life (holding the welfare of existing people equal) is given by the self-regarding value of that life *minus* α . (α is variously referred to as the ‘critical level’ (Blackorby, Bossert and Donaldson) or the ‘neutral level for existence’ (Broome)).

Total utilitarianism and critical level utilitarianism are structurally very similar (indeed, the former is a special case of the latter). Accordingly, if critical level utilitarianism is adopted, our reasoning proceeds much as above, with a simple modification. We obtain instead (at optimum population) that the self-regarding value of an average life exceeds the net negative other-regarding effect of that life by a quantity α , and that the difference in disvalue between a non-birth and a death is given by the sum of the transition factors plus the critical level α . However, advocates of critical level utilitarianism agree that the critical level cannot be too high (on pain of making the ‘negative repugnant conclusion’ too repugnant; see e.g. (Broome, 2004, pp. 213, 264)). It therefore seems likely that our main message, that other-regarding effects and transition effects are far more significant to the ‘value of a life’ than the dominant approach recognises — *unless* the world is significantly underpopulated — would survive the replacement of total utilitarianism with critical level utilitarianism.

We now take up this postponed issue of the possibility of underpopulation. We have noted that *if* the world is currently near or above its optimum population and if in addition premature deaths are bad, then other-regarding effects in general, and transition effects in particular, form (at least) a very significant part of the calculation of the (dis)value of a premature death. But it is worth noting well that any critical level utilitarianism, including total utilitarianism, will deem the world to be *underpopulated* in many cases in which untutored intuition would deem the world to be overpopulated. Intuitive thinking notes only that average welfare would be increased, or that some extant problems would be solved, by preventing some extra births; our theories recognise this, but recognise in addition that the lives thereby prevented are lives that would themselves have had positive value. We therefore do not conclude that other-regarding and transition effects *do* dominate the calculation of the disvalue of a premature death. An alternative way of resolving our tension is, guided by theory, to ditch the intuition that the world is currently overpopulated. If the world is *underpopulated*, premature deaths are indeed bad, for the simple reason that they reduce instantaneous population (an effect that can be entirely accounted for by factor (D1)), even if the deaths in question are of people who would not otherwise have gone on to have any additional children. The point I wish to press is that this — the underpopulation hypothesis — is the *only* way of reconciling the hypothesis that (D1) dominates the badness of a premature death with the above population axiologies discussed above.

Those axiologies, of course, are not uncontroversial; nothing in population axiology is uncontroversial.¹³ In particular, the tension between a high disvalue for prema-

¹³I have not had space to do justice here to anything like the full range range of possibilities, their relative merits and their respective implications for the present issue. The best account of the state of play in

ture deaths and a low disvalue for nonbirths could easily be avoided if the principle that Broome dubs ‘the neutrality intuition’ were true. This principle states that there is no value to adding an extra person to the population, except insofar as the addition modifies the well-being levels of those whose existence is held fixed. (In other words, the self-regarding value of *extra* lives has no ‘social’ value.) Against the backdrop of a population-axiological view incorporating this ‘neutrality’ principle, of course it is perfectly consistent to hold that the world is overpopulated — it would be bad to add an *extra* person — and simultaneously to hold that it is worth a very high price to preserve existing lives, on the grounds that the (positive) self-regarding value of the average life by far outweighs the (negative, and significant but relatively modest) net other-regarding value of that life. The intuitive force of the ‘neutrality’ principle is presumably responsible for the widespread failure to take seriously the possibility that other-regarding effects might be comparable in magnitude to the self-regarding disvalue of a premature death. As Broome has repeatedly emphasised, however (2005, 2004, 2012a), that principle, however intuitive, is extremely hard to defend.

A related issue is: What is the functional form of the badness of a death as a function of age at death? Any sensible exclusively self-regarding approach has the consequence that the badness of a death is a decreasing function of age of death: the earlier a person dies, the more *that person* loses by her premature death. This consequence is very plausible when the ages in question are more than (say) 25. But it is at least arguable that a death at age one year is worse than a neonatal death, and that a death at age 15 is worse than a death at age one year. Approaches that treat other-regarding effects as significant contributors to the overall social disvalue of premature deaths can accommodate these putative facts: if a child has survived to age 15 (for example) he has consumed large amounts of resources (via his sustenance and education) that he has not yet had time to pay back via working life, and the death of a fifteen-year-old will generally cause far more parental grief than the death of a newborn or a one-year-old. This is a further reason for thinking that other-regarding effects are non-negligible.

5 Conclusions

Pre-theoretically, it is intuitively highly plausible that the effect of a premature death on the person who dies dominates the calculation of the badness of that death. If that were the case, the WALY approach suggested by Broome would supply a viable method for working out, not only the *self-regarding* disvalue of a premature death, but also its *social* disvalue, i.e. its overall disvalue for the purpose of public policy analysis. However, unless the world is currently significantly underpopulated, that pre-theoretically plausible assumption does not remain plausible in the light of either the ‘total utilitarian’ axiology, or the ‘critical level utilitarianism’ that Broome has defended. If those axiologies are even serious contenders, then given approximately optimum or overpopulation, we are forced to take seriously the whole range of effects (D1)–(D3) and (T1)–(T5).

There are then two ways of proceeding. The first is directly to take on the task of quantifying each effect (D1)–(D3), (T1)–(T5) on a common scale, and summing them

population axiology is (Arrhenius, n.d.).

to obtain a figure for the monetary disvalue of a premature death. The slight advantage of this direct approach is that it enables us at least to avoid making any population-axiological (and therefore particularly controversial) assumptions in the special case in which future population size is independent of the timing of the death in question. Its disadvantage is that, barring *extremely* careful methodology, errors in the estimation of each factor are likely to swamp the calculation. The second, alternative, approach takes a detour via population axiology: we first make a judgement of which population axiology is correct and (in the light of that axiology) the value or disvalue of an extra birth, and we then modify that figure (e.g. by the sum of transition effects) to obtain a figure for the disvalue of a premature death. *Once the first step has been completed*, this calculation is relatively straightforward. This may well be a superior approach, since the first step must eventually be taken in any case. But there is no escaping the complexity: this second approach simply sweeps that complexity into the independent task of calculating the (dis)value of an additional birth. Given this complexity, it is hard to escape the conclusion that although precise calculations are possible *in principle*, in practice we are more likely to arrive at a sensible figure for the ‘monetary value of a life’ by sensible judgment than by precise calculation.

I will not attempt to venture a guess at a plausible figure. But if the world is currently optimally- or over-populated, it seems likely that the ‘value of a life’ figure that the arguments of this paper would sanction is significantly lower than the figures normally used. In that scenario, we have seen, the value of a life to the person living it tends to be cancelled or outweighed by the negative effects of that person’s resource consumption on others; any remaining ‘disvalue of a premature death’ may come only from transition effects; it is reasonable to expect these to be significantly lower.

It is of course innocuous to focus on calculations of the self-regarding disvalue of a premature death as *one input* into a sensible estimate of its social disvalue, always remembering that the result will need to be combined with other quantities before being put to use in practice. The literature suggests, however, that this is invariably *not* remembered in practice. In that case my conclusions are revisionary, not only of current practice in public policy analysis, but also of the practice that would ensue if, as Broome has suggested, WTP-based valuations of life were replaced with QALY/WALY-based calculations.

I will take the case of climate change as an illustrative example. It is often remarked — by Broome among others¹⁴ — that the deaths caused by climate change seem likely to account for a very significant fraction of the overall badness of that phenomenon. But the calculations backing up that claim involve using figures that are actually estimates of the *self-regarding* disvalue of a premature death as if they were estimates of the *social* disvalue of a premature death.¹⁵ Absent a case for thinking, not only that the estimates in question are at least roughly accurate *qua* estimates of self-regarding

¹⁴See Broome’s ((2000, p.956), (2012a, p. 156))

¹⁵The figures Broome cites in his (2000), cited also by the IPCC’s Second Assessment Report (IPCC, 1996, p.195), are from Fankhauser (Fankhauser, 1995). Fankhauser estimates total climate damages at US\$269.6 billion, of which \$49 billion, hence 18%, is due to mortality damages. The figure of \$49 billion is based on assuming ‘value of a life’ figures of \$1.5m for developed countries, \$300k for middle-income countries and \$100k for lower-income countries, these figures in turn being justified by the literature on willingness-to-pay estimates.

disvalue, but also that other-regarding contributions to overall (dis)value are negligible in comparison, we have precisely no reason for drawing any such conclusions. Of course the killing it does *may* be one of the worst things about climate change, but only in the sense that absent some credible analysis to guide our beliefs, anything may be the case.

Acknowledgements

I am grateful to John Broome and William MacAskill for comments on an earlier draft of this paper.

References

- Arrhenius, G. (n.d.). *Population ethics: the challenge of future generations*. (Forthcoming)
- Arthur, W. B. (1981, March). The economics of risks to life. *The American Economic Review*, 71(1), 54–64.
- Asian Development Bank. (2003). *Road safety guidelines for the Asian and Pacific region*. ADB Publishing.
- Blackorby, C., Bossert, W., & Donaldson, D. (1995, November). Intertemporal population ethics: Critical-level utilitarian principles. *Econometrica*, 63(6), 1303–1320.
- Blackorby, C., & Donaldson, D. (1984). Social criteria for evaluating population change. *Journal of Public Economics*, 25, 13–33.
- Broome, J. (n.d.). Measuring the burden of disease. In D. Wikler & C. J. L. Murray (Eds.), *Fairness and goodness: Ethical issues in healthcare resource allocation*. World Health Organisation. (Forthcoming)
- Broome, J. (1978). Trying to value a life. *Journal of public economics*, 9, 91–100.
- Broome, J. (1985, August). The economic value of a life. *Economica*, 52(207), 281–94.
- Broome, J. (1994). Structured and unstructured valuation. *Analyse & Kritik*, 16, 121–32.
- Broome, J. (1996, February). The value of life and the value of population. *Journal of population economics*, 9, 3–18.
- Broome, J. (2000, June). Cost-benefit analysis and population. *Journal of legal studies*, 29, 953–970.
- Broome, J. (2004). *Weighing lives*. Oxford University Press.
- Broome, J. (2005). Should we value population? *Journal of political philosophy*, 13, 399–413.
- Broome, J. (2012a). *Climate matters*. New York: W. W. Norton and Company.
- Broome, J. (2012b). *The philosophy and economics of the environment*. (Unpublished lectures delivered at the University of Oxford, Trinity Term 2012)
- Charles Blackorby, W. B., & Donaldson, D. (2005). *Population issues in social choice theory, welfare economics, and ethics*. Cambridge University Press.

- Cline, W. R. (1992). *The economics of global warming*. Washington, D.C.: Institute for International Economics.
- Fankhauser, S. (1995). *Valuing climate change: the economics of the greenhouse*. Earthscan.
- IPCC. (1996). *Climate change 1995: Economic and social dimensions of climate change: Contribution of working group III to the second assessment report of the Intergovernmental Panel on Climate Change*. Cambridge university press. (Edited by James P. Bruce, Hoesung Lee and Erik F. Haites)
- Mason, H., Jones-Lee, M., & Donaldson, C. (2009, August). Modelling the monetary value of a QALY: a new approach based on UK data. *Health economics*, 18, 933–50.
- Mishan, E. (1971). Evaluation of life and limb: a theoretical approach. *Journal of Political Economy*, 79, 687–705.
- National Institute for Health and Clinical Excellence. (2008). *Social value judgments: Principles for the development of nice guidance*. NICE. (Second edition)
- Nord, E. (1992, March). Methods for quality adjustment of life years. *Social Science and Medicine*, 34, 559–69.