Climatic Change (2015) 130:347–358
DOI 10.1007/s10584-014-1131-2

Casualties as a moral measure of climate change

John Nolt

Received: 20 January 2014 / Accepted: 13 April 2014 / Published online: 26 April 2014

Abstract Climate change will cause large numbers of casualties, perhaps extending over thousands of years. Casualties have a clear moral significance that economic and other technical measures of harm tend to mask. They are, moreover, universally understood, whereas other measures of harm are not. Therefore, the harms of climate change should regularly be expressed in terms of casualties by such agencies such as IPCC's Working Group III, in addition to whatever other measures are used. Casualty estimates should, furthermore, be used to derive estimates of casualties per emission source up to a given date. Such estimates would have wide margins of error, but they would add substantially to humanity's grasp of the moral costs of particular greenhouse gas emissions.

A recent commentary in *Nature* contended that methane releases from permafrost beneath the East Siberian Sea will cost the global economy $60 trillion (Whiteman et al. 2013). If true, this is, I suppose, really bad. (I say "if true" because the estimate is controversial (Mooney 2013).) And yet this article, like many others with similar monetary conclusions, obscures a fundamental ethical point: climate change does not just cost money; it causes large numbers of people to suffer and die.

Monetary estimates of the damages of climate change are, of course, useful. Cost estimates for Arctic methane, for example, should, as the authors of the *Nature* commentary suggest, be continually refined and accounted for, as they are not today, in such projections as the World Economic Forum's *Global Risk Report* and the International Monetary Fund's *World Economic Outlook.* But far too much of the literature on the harms of climate change is formulated exclusively in economic terms.

Here I argue that casualty estimates are among the most morally significant measures of the harm of climate change, that they should figure prominently in the ethical considerations of such advisory bodies as Working Group III of the Intergovernmental Panel on Climate Change (IPCC), and that they should be widely disseminated. I maintain, further, that they can be used to calculate another measure, casualties per emission source to a given future date, that provides a sense of the moral significance of greenhouse gas (GHG) emissions by nation, region, industry and perhaps even per person.

This article is part of a special issue on "Multidisciplinary perspectives on climate ethics" with guest editors Marco Grasso and Ezra M. Markowitz.

J. Nolt (✉)
Department of Philosophy, University of Tennessee, Knoxville, TN 39776-0480, USA
e-mail: nolt@utk.edu

1 The scope of the harm

One difficulty in communicating the dangers of climate change to policymakers and the public lies in explaining the magnitude of the harm. Climate change takes place globally, over time spans of unfamiliar magnitude. The IPCC fourth assessment report predicts that "climate change over the next century is *likely* to adversely affect hundreds of millions of people through increased coastal flooding, reductions in water supplies, increased malnutrition and increased health impacts" (IPCC 2007: 65). But that will not be the end of it. Recent results that incorporate slow climate feedbacks due to Earth system sensitivity suggest that given the likely climate sensitivity (3 °C for doubling of CO_2) anthropogenic warming will persist for 23,000 to 165,000 years (Zeebe 2013). Hence people may continue to be harmed by today's emissions for an extraordinarily long time.

Biodiversity loss is an even longer-term prospect. Many biologists worry that climate change together with other human impacts will eventually precipitate a mass extinction—that is, loss of most of the world's species in a geologically brief time. Barnosky et al., point out that recovery of biodiversity from such an event "will not occur on any timeframe meaningful to people: evolution of new species typically takes at least hundreds of thousands of years and recovery from mass extinction episodes probably occurs on timescales encompassing millions of years" (Barnosky et al. 2011: 51). I have discussed the ethical implications of long-term biodiversity loss elsewhere (Nolt 2011a). This paper focuses on harms to humans.

Because climate change does not end at any definite future date, meaningful numbers for aggregate harms must be indexed to a definite time period—a particular year, for example, or decade or century or millennium. But in the absence of an explanation that the harm may continue thereafter, the statistic is misleading. Thus the IPCC's mention of those hundreds of millions of people who will likely be adversely affected over the next century, startling though it is, obscures the fact that worse may occur in centuries following. We need a sense, too, of this bigger picture.

2 Ethical issues with economic cost estimates

We will not get it from economic cost estimates. Such estimates, if they account for future death and suffering, typically employ discounting procedures that devalue them. Consider, once again, the study described in the *Nature* commentary, which uses fairly standard methodologies to estimate economic costs for methane releases. Authors Gail Whiteman, Chris Hope and Peter Wadhams obtained their estimates by running the PAGE09 integrated assessment model 10,000 times under various assumptions.

In PAGE09 non-economic impacts, including those to human health, are expressed in monetary units as a polynomial function of regional temperature. These costs are then discounted by a pure time preference rate—an annual percentage by which all future values (whether costs or benefits) are reduced toward zero (Hope 2011). Use of such a discount rate is standard in economic assessments. Historically, it has been justified by the claim that that our preferences for future goods or avoidance of future evils weaken the further they are expected to lie in the future.

But any such rate is difficult to justify empirically. In a comprehensive survey of the history of the practice, Frederick, Loewenstein, and O'Donoghue conclude that

> … virtually every core and ancillary assumption of the DU [discounted utility] model has been called into question by empirical evidence collected in the past two decades. …

> While the DU model assumes that intertemporal preferences can be characterized by a single discount rate, the large empirical literature devoted to measuring discount rates has failed to establish any stable estimate (Frederick et al. 2002: 393).

Yet it is precisely such a single, stable discount rate that is assumed in integrated assessment models, such as PAGE09 and in many similar economic analyses.

More importantly, pure time preference discounting is, at least as regards matters of life and health, ethically indefensible. It considers the preferences of future people only insofar as they matter to us at present, and so unjustly discriminates against those future people. (Imagine, by way of analogy, taking the preferences of one race or gender into account only insofar as they mattered to members of another race or gender.) Hence we should reject pure time discounting. (Philosopher and economist John Broome (2012: 148–53) reaches this same conclusion, though by a different route.)

A common rejoinder is that we *can't* take the preferences of future people into account since we have no way of knowing their preferences. But that is false. Human preferences for life and health remain relatively constant over time. In matters of life and health, therefore, we should not subject loss of life or health to pure time discounting, which even at modest discount rates reduces the values of the lives and health of future people to a tiny fraction of the values of ours within a century or so.

Economists sometimes give other justifications for pure time discounting. In his work on climate economics Nicholas Stern adopts a pure time annual discount rate of 0.1 %, which reflects, not the human psychological tendency to discount future harms and benefits, but what he takes to be the annual probability of human extinction (Stern 2008). Unlike that human psychological tendency, the probability of human extinction is clearly morally relevant. If an asteroid impact will eliminate humans by 2050, for example, then, so far as humans are concerned, we need not worry about the effects of climate change thereafter. But in evaluating the harm of various scenarios, clarity is best served by separating the probability of the scenario from the harm, rather than treating probability as a factor in harm.

Suppose, to take an unrealistically simple example, we wish to estimate weather-related economic losses in 2050 on just two different GHG emissions scenarios. Let v_1 and v_2 be the loss estimates, respectively, for scenario 1 and scenario 2. And let p be the probability (the same for each scenario, let's assume) of human extinction by 2050. If we follow Stern, then, we would, for example, express the loss in scenario 1 as $v_1 - v_1 p$ (and similarly for scenario 2). But this makes it easy to overlook the fact that if humanity survives to 2050 and scenario 1 occurs, the loss actually experienced will be v_1.

For ethical purposes at least, it is more informative to forego the discounting and distinguish two potential outcomes for each scenario: the outcome if humanity survives to 2050 and the outcome if not, thus clearly distinguishing actual from expected loss. The actual loss under scenario 1 is v_1 if we survive, and (presumably) 0 if we don't, since in that case we are already dead. Probability-based discounting yields only the *expected* loss for a given scenario—that is, the sum of the probability-weighted values for each of that scenario's possible outcomes. The expected loss in scenario 1, for example is $v_1 \times (1-p) + 0 \times p = v_1 - v_1 p$, the sum of probability-weighted losses for the two potential outcomes of scenario 1.

The advantage for ethics, which is concerned with outcomes that we can control, of this procedure is that it clearly separates losses due to the emissions scenarios that we may choose from those due to other contingencies (e.g. asteroid impacts) over which we may have no control. (If we *can* prevent asteroid impacts, then that is all the more reason treat them as separate outcomes, so that options for preventing them, together with their costs, are explicitly considered.)

Various other justifications are given for discount rates, but none are morally defensible when applied to the harms of bodily injury and death. (For a critique of these other justifications, see Parfit (2004), Appendix F.)

Preferences matter in neoclassical economics—and hence show up in applications like PAGE09—because the ultimate aim of that discipline is preference satisfaction, under the assumption that preference satisfaction constitutes human welfare. Many economists ascertain preference strength by people's willingness to pay. So on their view, the value of anything (including life and health) is determined by people's willingness to pay for it.

Two of the assumptions on which this view is based are false. First, willingness to pay is not an adequate measure of preference strength, for it reflects not only preference strength but disposable wealth. Hence, unless corrected for disposable wealth, it is at best a measure of wealth-weighted preference strength—a measure, in other words, that is biased in favor of the rich.

Second, preference satisfaction is not the same thing as welfare. There are a number of reasons for this, including all of the following:

- Consumer preferences (which are the ones primarily reflected in economic cost estimates) are distinct from considered values, whose satisfaction may more closely reflect actual welfare (Sagoff 1988).
- Preferences are frequently manipulated by propaganda and advertising so that they have little to do the welfare of the individuals whose preferences they are.
- Some preferences (e.g., for fatty or sugary food) are biological in origin and dysfunctional under conditions of abundance, so that their maximal satisfaction is objectively harmful.
- Some preferences are the result of addictions; again satisfying them is objectively harmful.
- Often people are benefited by events for which they had no preferences.
- Sometimes people are benefited more by eliminating than by satisfying certain preferences (e.g., preferences for conspicuous consumption).

Preference satisfaction and human welfare are thus distinct. This has led some theorists who assume that the aim of policy is to promote human welfare to seek a more accurate conception of welfare.

3 Objective conceptions of welfare and harm

Ethical theorists who reject the preference-satisfaction conception of welfare generally opt either for a considered preference view, on which welfare is determined only by informed and rationally considered preferences, or for an objective welfare view (sometimes called an "objective list view") on which human welfare consists in the satisfaction certain objective conditions for a decent human life. These may include, for example, health, adequate wealth, longevity, security, freedom from oppression, education, or the like.

But considered preference satisfaction is hardly an adequate proxy for human welfare either. There are, for example, apparently rational people, as fully informed and cognizant as anyone, who nevertheless prefer to play video games all day, eat self-destructively, smoke, or use cocaine. One can object that these people are not *perfectly* rational or *perfectly* informed. But the problem is that there are no perfectly rational and perfectly informed people. Hence if we have some idea of what such people would prefer, we certainly did not get it from them. It is, more likely, just our best guess as to what is objectively good for people. In any case, the notion of considered preference satisfaction has not to my knowledge yielded any conception of welfare that would be practical for assessing changes in welfare (whether harms or benefits) resulting from climate change. So we may set it aside.

Objective welfare conceptions—which bypass the question of what we subjectively prefer, or rationally ought to prefer, entirely—are more promising. One fairly objective measure that is already used in some policy applications is the Human Development Index (HDI) (UNDP 2013). Nominally a measure of the level of development in a region, the HDI, which combines longevity with measures of education and living standards, can also be regarded as a rough indicator of average human welfare for that region.

Another objective welfare measure, commonly used in the field of public health, is the disability-adjusted life year (DALY). One DALY is one year of healthy life. Harm can be expressed as loss of DALYs relative to statistical life expectancy. Thus (without further modifications; see below) the harm of death to a person with a life expectancy of 80 years who dies at age 30 is 50 DALYs. DALYs also account for disability resulting from injury or disease as reduction in the value of a year lived with the disability. If, for example, the 30-year-old is disabled instead of killed, but lives 50 more years with the disability, the loss is some fraction of 50 DALYs, the fraction depending on the severity of the disability. More generally, each disability is assigned a value between 0 (death) and 1 (health), lower numbers representing greater severity. A disability with value 0.4, for example, reduces life to 40 % of its healthy value. So, for instance, if an individual endures for 50 years a disability with value 0.4, then the loss is $0.6\times50=30$ DALYs.

The DALY loss for a population resulting from a given cause (a hurricane, for example, or malnutrition in a given year) is an aggregate measure of the harm to members of that population from that cause.

Various modifications are sometimes made to DALY measures. DALYs used in some World Health Organization reports, for example, give less weight to years lived at younger and at older ages than to years in the middle of life. Some also include a fixed annual discount rate, like the economic rate discussed above, that expresses the so-called “social preference” for goods that occur sooner rather than later (WHO 2013). But because such weightings are controversial, the WHO often provides statistics both with and without them.

Lost DALYs are, in some respects, a more objective measure of harm than are economic costs. But non-empirical assumptions are needed to obtain numerical results. Just how many lost years of life are equivalent, for example, to living for *x* years with the debilitating effects of emphysema? To what degree are the “magic” years of childhood and the “golden” years of old age less valuable than the middle years of adulthood? To what degree are years of life less valuable as a result of occurring later in time? Is a year of healthy life at a given age worth the same for everyone? These are not strictly empirical questions, yet DALY totals depend on mathematically precise answers to them. By changing these answers, one can re-order the rankings of harms expressed as DALYs (Arnesen and Kapiriri 2004).

DALYs have, in addition, one great practical disadvantage. The public does not understand them. If to the question, “how much harm will climate change do in the next decade (or century)” one answers with a certain large number of DALYs lost, this will mean nothing to most people. If, on the other hand, one answers with a casualty or fatality estimate, nearly everyone will understand.

4 Existing casualty and mortality estimates

Casualty estimates are another objective measure of harm, but they are much simpler and more readily comprehensible, though perhaps less informative, than DALYs. By a *casualty estimate*

I mean an estimate of how many people die or suffer injury or illness from a given cause over a specified period of time. *Mortality* or *fatality estimates* (or, more colloquially, *death tolls*) are similarly defined, but they consider only deaths, not injuries or illnesses.

The IPCC statement that hundreds of millions of people will suffer adverse effects from climate change over the next century is something like a casualty estimate. The IPCC considers four kinds of adverse effects: increased coastal flooding, reductions in water supplies, increased malnutrition and increased health impacts. The adverse effects of coastal flooding include, presumably, economic losses and forced migration. Reductions in water supplies and increased malnutrition create thirst and hunger. Among health impacts the IPCC lists:

> … increases in malnutrition; increased deaths, diseases and injury due to extreme weather events; increased burden of diarrhoeal diseases; increased frequency of cardio-respiratory diseases due to higher concentrations of ground-level ozone in urban areas related to climate change; and the altered spatial distribution of some infectious diseases (IPCC 2007: 48).

The IPCC's characterization is redundant; malnutrition, for example, is a health effect and so should not be listed *in addition* to health effects. Moreover it includes some harms (e.g., forced migration) that are probably not best characterized as casualties. Meaningful casualty estimates require an agreed-upon definition of "casualty." Still, the IPCC statement provides some idea of what a long-term casualty estimate for climate change would look like.

There have been at least three efforts to quantify mortality rates for climate change. The first was by the World Health Organization, which in 2009 estimated that climate change had been causing about 141,000 excess deaths annually by the year 2004 due to disease and malnutrition. This figure was inferred from public health data by comparing observed levels of malnutrition, diarrhea, malaria and dengue fever with their levels in a counterfactual situation representing average climate conditions during the period 1961–1990, when the effects of anthropogenic climate change were presumably negligible. It did not, however, include deaths due directly to high temperatures, flooding, or the increased severity or frequency of storms. The WHO also provided DALY losses for the same data (WHO 2009 Table A3, p. 50).

The second attempt—a 2009 study by the nonprofit Global Humanitarian Forum—used the WHO data but also incorporated estimates of deaths from storms and weather-related flooding. It estimated the death toll from climate change at 300,000 people annually. Mortality rates for storms and flooding were obtained by comparing reported trends in the frequency of weather-related disasters to those for geophysical disasters, such as earthquakes. The latter were assumed to be unaffected by climate change. The period considered was 1980–2005. While numbers of loss-producing disasters for both weather-related and geophysical events increased—due, perhaps, to better reporting and population increases—those for weather-related disasters increased by about 40 % more. The report therefore attributed 40 % of weather-related deaths to climate change. This same 40 % figure was also used in estimating economic losses and numbers of people seriously affected by climate change (GHF 2009).

The third attempt, 2012 report by Development Assistance Research Associates (DARA 2012), considered data for many specific factors (including drought, floods, landslides, storms, wildfires, diarrhea, heat and cold illnesses, hunger, malaria, and vector-borne meningitis). It estimated the current annual death toll from climate change at around 400,000. Assuming business as usual, it predicted, moreover, that by 2030 that number would rise to nearly 700,000. Extrapolating from this third study, Broome has concluded that climate change will cause tens of millions of deaths by the end of this century (Broome 2012, 33).

No doubt the methodologies of these studies are imperfect. Yet they demonstrate that statistics and methods for providing rough casualty estimates and projections are already available. There is, moreover, reasonably close agreement among existing estimates (though these are not wholly independent of one another). Clearly there is need for critique and further refinement of such estimates.

5 Moral importance of casualty estimates

One immediate advantage of casualty and fatality numbers is that they are indisputably empirical. When (as for many past or present events) actual counts can be made, the results are hard numbers, purely matters of fact. Where estimates are needed, they are still estimates of what are or will be hard facts.

A second advantage is that casualty and fatality estimates are familiar to everyone and much better understood by the public than any of the more technical measures introduced by economists or public health experts. The news media report them daily for all sorts of calamitous events. Historians rightly use them to gauge the relative severities of disasters (battles, storms, earthquakes, plagues, etc.) across wide stretches of time.

A third advantage is that casualty and fatality estimates may be more likely than some other ways of assessing the dangers of climate change to elicit constructive public concern. A recent study comparing audience reactions to articles framed in terms of dangers to the environment, national security, or public health, found that

> … across audience segments, the public health focus was the most likely to elicit emotional reactions consistent with support for climate change mitigation and adaptation. Findings also indicated that the national security frame may possibly boomerang among audience segments already doubtful or dismissive of the issue, eliciting unintended feelings of anger (Myers et al. 2012).

Death, disease and injury are, of course, matters of public health.

Casualty estimates are typically divided into fatalities and injuries—which, let us assume, are defined broadly enough to include instances of serious illness. (In military reports, soldiers captured are often included as well, but capture is irrelevant here.) This division is more informative than simply adding fatalities and injuries together, and it is methodologically appropriate, because death and injury are not directly comparable harms. One can put them on the same scale, but that requires deciding how much moral value to give injuries of different types—a decision that is to some degree arbitrary, non-empirical, and controversial. Casualty numbers, partitioned into dead and injured, however, require no such weightings. For past or present casualties, they are just raw empirical data. Hence this partitioning is generally best.

Casualty estimates do not, of course, provide a complete account of harm. They do not, like DALYs, take into account the duration or severity of disabilities or consider the number of years lost per death. Both DALYs and casualty estimates, moreover, ignore other sorts of harm, such as economic losses.

In addition, both casualty estimates and DALYs are insensitive to the fact that climate change produces some benefits, as well as costs, so neither gives us a sense of the balance of costs and benefits of a given policy. Large numbers of casualties would be associated, too, with precipitous reductions in the use of fossil fuels, if no replacement were provided.

How important these omissions are depends on one’s ethical framework. Policy decisions are generally made within either of two broad frameworks: deontology and consequentialism.

Deontological ethics is rule-based, justice-oriented, and often expressed in the language of human rights. Deontologists evaluate actions by their accord with rules, not by their consequences. They generally regard proper treatment of individuals as more important than the welfare of the whole. Vanderheiden (2008) develops such a view in some detail. Whether our emissions produce greater benefits for us, or even for both us and future people, is on most deontological views irrelevant. The fact that they produce innocent casualties is sufficient to condemn them morally. Casualty estimates are thus directly relevant to deontologically justified policy decisions, even if not accompanied by an estimate of the economic benefits of climate emissions.

Much policy, however, is formulated within a consequentialist framework. For consequentialists policy should depend on anticipated aggregate consequences. Casualties are also relevant to consequentialist decision-making, but they must be aggregated with all other welfare losses or gains. Ideally this is done by measuring all consequences on the same numerical scale—for example, the Human Development Index. Casualties are accounted for by the HDI, since one of its components is longevity. All such single-scale measures must, however, employ non-empirical weighting schemes to combine the relevant factors (longevity, education and wealth in the case of the HDI) into a single scale.

One can avoid the need for such a weighting scheme by refusing to combine the various scales. Multiple scales have, however, the disadvantage of complicating decision procedures.

My point is not to propose any particular objective welfare measure or measures, but merely to note that any reasonable objective welfare consequentialist decision procedure will in one way or another use something like casualty estimates as an input. Casualty estimates are, in other words, significant no matter which ethical framework we use.

I have been touting the advantages casualty estimates and decrying the disadvantages of economic assessments and public health assessments quantified as DALYs, yet my point is not that casualty estimates should replace either of these. The point is merely that in assessing the harms of climate change, analysts should in addition, and separately, provide casualty estimates as a moral measure of climate change.

This raises the question of whether shadow prices for injuries or fatalities (as opposed to their actual market costs) should then be removed from economic analyses, so that economic costs and bodily harms to humans are disentangled. Since that would avoid potential double counting and would simplify and reduce some of the arbitrariness of economic analyses, it seems worth considering, especially since there is no reason why economic analyses could not continue to be done in the current ways as well. Information should be made available to policymakers in whatever format they find illuminating.

6 Duration of harms

Current annual casualty estimates for climate change are considerably lower than those for other important causes of death and injury. Annual death toll estimates for anthropogenic air pollution, for example, are in the low millions (Silva et al. 2013), roughly an order of magnitude greater than the annual death toll from climate change. This might lead us to regard casualties from climate change as relatively insignificant. But that would be a mistake.

The reason, in a word, is: duration. For most forms of anthropocentric air pollution (particulates, ozone, sulfur dioxide, nitrogen oxides, industrial chemicals), once emissions cease, the pollutants are rapidly dispersed, diluted or degraded, so that health effects, which were predominantly regional, diminish quite rapidly. But the harms of emissions of carbon dioxide, the most important contributor to climate change, are, as was noted at the outset, global and much longer-lasting.

In part this is because the warming effects of carbon dioxide are mechanical, not chemical. Most other pollutants cause harm via chemical reactions that transform them into something else and thus terminate their effect. But CO_2 raises temperatures by absorbing infrared radiation from the earth's surface and dispersing the resulting heat into the environment. CO_2 molecules are not chemically transformed in the process. They therefore trap excess heat for as long as CO2 levels remain elevated—which is on the order of centuries (Archer, et al. 2009). How hot it gets will be determined by the total quantity of CO_2 we emit. More specifically, maximum warming over the coming centuries will, relative to pre-industrial times, be nearly proportional to cumulative total anthropogenic carbon emissions (Stocker 2013). Once all emissions cease, there will be little further temperature increase; but once temperatures peak, they will (due to slow climate feedback mechanisms) decline only slowly, remaining elevated for thousands of years (Matthews and Solomon 2013; Zeebe 2013). Since the effects of global warming are already on balance harmful, high temperatures over these millennia will, presumably, contribute to additional harms by way of storms, floods, drought, and so on. Adaptation efforts may, of course, diminish the rate of harm, but only by imposing additional economic or human costs.

This year's carbon dioxide emissions will, consequently, contribute to harms occurring not only over the coming decades, but probably over thousands of years. Annual mortality rates, and even the tens of millions of deaths that Broome projects for this century, are only the tip of the iceberg.

7 Casualties per emission source

We could get a more specific sense of the moral significance of our actions if we could estimate casualties per emission source—e.g., nation, region, industry, perhaps even individual (Nolt 2011b and 2013)—up to a given time. To do this, we need to know how casualties depend upon emissions. We have already noted that maximum temperature is roughly proportional to cumulative total emissions and that after the maximum is reached, temperature will decline only slowly. How, then, do harms depend on temperature?

Other things being equal, harms from climate change over a given time period depend directly on the elevation in global average temperature. ("Other things being equal" means that we ignore such confounding factors as population size and adaptation measures.) To put it simply, a hotter year is a more deadly year. This is, to the best of my knowledge, uncontroversial. PAGE09, for example, as was noted above, assumes that harm, both economic and non-economic, is a polynomial function of temperature elevation.

Given that maximal temperature elevation depends on cumulative emissions and harm depends on temperature, it follows that, other things being equal, in the long run harm depends on cumulative emissions. We can therefore to a first approximation simplify the problem of attributing casualties (and other forms of harm as well) to a given emitter by assuming that:

(A) The harm attributable during some period—say, year y— to an earlier emission of c metric tons of carbon dioxide is the total harm from anthropogenic climate change during y multiplied by c/a, where a is the cumulative total anthropogenic emission up to y.

We may think of cumulative emissions as the sum of annual emissions, starting with some year, say 1850, before which emissions were negligible.

This assumption is inaccurate three respects. First, the proportionality between cumulative total emissions and maximal temperature is only approximate. Second, the proportionality between cumulative total emissions and temperature may not hold prior to the achievement of maximal temperature. And, third, the earlier the emission of *c* metric tons occurs, the more warming it will have accomplished, since it will have been trapping heat and triggering feedback mechanisms for a longer time. A full account of how much an emission contributes to temperature would have to adjust for all three inaccuracies, but I will not attempt that here. In any case, the latter two diminish with the length of the time period considered (especially as it stretches beyond the fossil fuel era) and hence eventually become insignificant.

Given assumption (A) and a record of the emissions of a given emitter, we can then, given one additional set of data, estimate the harm attributable to that emitter. These additional data are estimates of total annual harm from climate change, expressed, for example, in DALYs or lost dollars or casualties or deaths. We saw examples of such estimates above. Plotting time on the *x*-axis and total harm year-by-year on the *y*-axis, we can then fit a continuous curve to these annual estimates. The curve represents rate of harm, which might be expressed in such units as casualties per year. Thus for any time *t* as far out as our estimates go, we obtain an instantaneous rate *r*(*t*) of harm from climate change at *t*.

Now for any GHG emission source *s*, let *ce*(*s*,*t*) be the cumulative emissions of *s* from 1850 to *t*, expressed, for example, in gigatons of carbon.

According to assumption (A), the rate of harm *r*(*s*,*t*) caused by the emissions of *s* at *t* is

$$r(s,t) = r(t) \times \frac{ce(s,t)}{ce(a,t)}$$

where source *a* is all anthropogenic emitters. (Note that the rate of harm for *s* at *t* is in general nonzero even if the emissions of *s* have ceased, since the temperature contributions of its emissions continue.) Then total harm *h*(*s*,*t*) caused by *s* up to time *t* is, again to a first approximation,

$$h(s,t) = \int_{1850}^{t} r(s,x)dx.$$

In theory, given the long duration of temperature elevation, we could obtain the totality of the harms of *s*, by setting *t* to something on the order of 100,000. But of course such estimates are wholly impractical.

In practice, however, we already have estimates (from the DARA report) that extend *t* to 2030, and we know the emission histories of many sources. Thus shorter-term estimates of casualties or death tolls attributable to a given emitter are already practical. These may help to give emitters a robust moral sense of what is at stake in their emissions.

8 Conclusions

Climate change does not just cost money; it causes large numbers of casualties over thousands of years. Casualties have a moral significance that economic and technical measures of harm tend to mask. They are, moreover, universally understood, whereas other measures of harm are not. Therefore, the harms of climate change should be expressed in terms of casualties by such agencies such as IPCC's Working Group III, in addition to whatever other measures are used.

This would require adoption of some standard of what counts as a casualty—ideally one that facilitates comparison with fatal and injurious events throughout history.

Casualty estimates should, furthermore, be used to derive estimates of casualties per emission source up to times as far into the future as possible. Such estimates would have wide margins of error, but they would add substantially to humanity's grasp of the moral costs of specific GHG emissions.

Acknowledgments I wish to thank Jenna Nolt and Robert Nowell for help with the research for this paper and Klaus Keller for helpful criticisms. Three anonymous reviewers for this journal provided astute suggestions that have substantially improved this paper

References

Archer D et al (2009) Atmospheric lifetime of fossil fuel carbon dioxide. Annu Rev Earth Planet Sci 37:117–134. doi:10.1146/annurev.earth.031208.100206
Arnesen T, Kapiriri L (2004) Can the value choices in DALYs influence global priority-setting? Health Policy 70: 137–149. doi:10.1016/j.healthpol.2003.08.004
Barnosky A et al (2011) Has the earth's sixth mass extinction already arrived? Nature 471:51–57. doi:10.1038/nature09678
Broome J (2012) Climate matters: ethics in a warming world. W. W. Norton, New York
Development Assistance Research Associates (DARA) (2012) Climate vulnerability monitor 2nd ed. http://daraint.org/climate-vulnerability-monitor/climate-vulnerability-monitor-2012/report/. Accessed 6 November 2103
Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: a critical review. J Econ Lit XL:351–401
Global Humanitarian Forum (GHF) (2009) Climate change: the anatomy of a silent crisis. http://www.ghf-ge.org/human-impact-report.pdf. Accessed 6 November 2103
Hope, C (2011) The social cost of CO_2 from the PAGE09 model, Economics Discussion Papers 2011–39, http://econstor.eu/bitstream/10419/49936/1/668556455.pdf. Accessed 6 November 2103
Intergovernmental Panel on Climate Change (IPCC) Climate Change (2007) Synthesis Report. Cambridge University Press, Cambridge
Matthews H, Solomon S (2013) Irreversible does not mean unavoidable. Science 340:438–439. doi:10.1126/science.1236372
Mooney, C (2013) How much should you worry about an arctic methane bomb? Mother Jones, 8 August, 2013, http://www.motherjones.com/environment/2013/08/arctic-methane-hydrate-catastrophe?page=1. Accessed 6 November 2103
Myers T, Nisbet M, Maibach E, Leiserowitz A (2012) A public health frame arouses hopeful emotions about climate change. Clim Chang 113:1105–1112. doi:10.1007/s10584-012-0513-6
Nolt J (2011a) Nonanthropocentric climate ethics. WIRES Climate Change 2:701–711. doi:10.1002/wcc.131
Nolt J (2011b) How harmful are the average American's greenhouse gas emissions? Ethics, Policy Environ 14, 1: 3–10. doi:0.1080/21550085.2011.561584
Nolt J (2013) Replies to critics of 'How harmful are the average American's greenhouse gas emissions?' Ethics, Policy Environ 16(1):111–119. doi:10.1080/21550085.2013.768399
Parfit D (2004) Reasons and Persons. Clarendon Press, Oxford
Sagoff M (1988) The economy of the earth: philosophy, law and the environment. Cambridge University Press
Silva R et al (2013) Global premature mortality due to anthropogenic outdoor air pollution and the contribution of past climate change. Environ Res Lett 8:3. doi:10.1088/1748-9326/8/3/034005
Stern N (2008) The economics of climate change. American Econ Rev 98(2):1–37. doi:10.1257/aer.98.2.1
Stocker T (2013) The closing door of climate targets. Science 339:280–282. doi:10.1126/science.1232468
United Nations Development Programme (UNDP) (2013) 2013 Human Development Report. Human development index. http://hdr.undp.org/en/statistics/hdi/. Accessed 6 November 2103
Vanderheiden, S (2008) Atmospheric justice: a political theory of climate change. Oxford University Press
Whiteman G, Hope S, Wadhams P (2013) Vast costs of arctic change. Nature 499:401–403. doi:10.1038/499401a

World Health Organization (WHO) (2009) Global health risks: mortality and burden of disease attributable to selected major risks. http://www.who.int/healthinfo/global_burden_disease/GlobalHealthRisks_report_full.pdf . Accessed 6 November 2103

World Health Organization (WHO) (2013) Disability weights, discounting and age weighting of DALYs. http://www.who.int/healthinfo/global_burden_disease/daly_disability_weight/en/index.html. Accessed 6 November 2103

Zeebe R (2013) Time-dependent climate sensitivity and the legacy of anthropogenic greenhouse gas emissions. Proc. Nat. Acad. Sci 110. doi:10.1073/pnas.1222843110