

[Draft – not for citation or quotation without permission]

**Regulation, Compensation, and the Loss of Life:  
What Cost-Benefit Analysis Really Requires**

**Ty Raterman**

Department of Philosophy  
University of the Pacific  
3601 Pacific Ave.  
Stockton, CA 95211

[traterman@pacific.edu](mailto:traterman@pacific.edu)

(209) 946-7624

## **Regulation, Compensation, and the Loss of Life: What Cost-Benefit Analysis Really Requires**

### **Abstract**

This paper explores the question of whether cost-benefit analysis (CBA) will ever endorse what I term “lethal policies,” which are roughly those that are expected to result in more human deaths than would another available policy. Both supporters and opponents of CBA commonly maintain that it will frequently do so. I argue, however, that this is a mistake. I first show that many of the leading contemporary teachers of and commentators on CBA mischaracterize it as demanding only that a policy’s expected benefits outweigh its expected costs. What CBA really requires is that it be possible for gainers to compensate losers, which is different in ways on which I elaborate. In addition to being unfaithful to what those who pioneered CBA took as critical to the justification for policy changes, these contemporary mischaracterizations also obscure the difficulties that CBA has in endorsing lethal policies. I discuss how economists who are sensitive to the demand for compensability argue that this can be satisfied even where those who lose through a policy lose so badly as to die. After explaining why their argument is unsatisfying, I advance a conception of costs and benefits that is superior to the standard conception and that clearly entails that the costs imposed by death are (in many cases) non-compensable and thus that CBA will not (except in rare cases) approve of lethal policies.

# **Regulation, Compensation, and the Loss of Life: What Cost-Benefit Analysis Really Requires\***

## **I. Introduction**

There is considerable disagreement regarding the merits of cost-benefit analysis (henceforth, CBA) as a tool for appraising public policy proposals. To some, it is the paragon of rational decision making.<sup>1</sup> For others, it is not only besieged by technical problems but also plainly morally indefensible.<sup>2</sup> Many of the moral criticisms pertain to how CBA handles those public policy decisions that have direct implications for human mortality. This especially includes decisions regarding environmental regulatory standards, since an enormous number of deaths are attributable to what can broadly be construed as environmental factors—especially organic and chemical pollutants. Research has put the number at as high as roughly 40% of all worldwide deaths.<sup>3</sup> Consider, for example, the World Bank estimates that on an annual basis in China, 350,000 to 400,000 people die prematurely from outdoor air pollution, roughly 300,000 die from indoor air pollution, and roughly 60,000 from water pollution.<sup>4</sup> In the United States: (1) the EPA suggests that by the year 2020, the Clean Skies Act will prevent between 8,400 and

---

\* Thanks go to Douglas MacLean for help in the development of the core ideas in this paper, and more generally for all the time he has spent with me discussing cost-benefit analysis, nonmarket valuation, and the like. Thanks are also due to Myrick Freeman, Ramona Ilea, and Peter J. Meyer for insightful feedback.

<sup>1</sup> For a thoughtful and readable case for CBA, see Cass Sunstein, *The Cost-Benefit State: The Future of Regulatory Protection* (Chicago: American Bar Association, 2002).

<sup>2</sup> For a pithy and provocative critique of CBA, see Frank Ackerman and Lisa Heinzerling, *Priceless: On Knowing the Value of Everything and the Price of Nothing*. (New York: The New Press, 2004).

<sup>3</sup> David Pimentel *et al.*, “Ecology of Increasing Disease: Population Growth and Environmental Degradation,” *BioScience* 48 (1998), pp. 817-826. Also, Pimental *et al.*, “Ecology of Increasing Diseases: Population Growth and Environmental Degradation,” *Human Ecology* 35 (2007) (forthcoming).

<sup>4</sup> Joseph Kahn and Jim Yardley, “As China Roars, Pollution Reaches Deadly Extremes,” *New York Times* (August 26, 2007), p. A1.

14,000 premature deaths annually;<sup>5</sup> (2) the mandate of smoke-free work environments in Florida prevents approximately 1,800 premature deaths each year;<sup>6</sup> (3) undoing current OSHA regulations governing methylene chloride would lead to approximately nine more workplace deaths in the United States a year.<sup>7</sup> Examples of this sort are plentiful. As a shorthand, let us call a policy that would reasonably quickly lead either to deaths when another available policy would not, or to more deaths than would another available policy, a “lethal policy.”<sup>8</sup> This paper explores complexities surrounding whether CBA is capable of endorsing lethal policies. Both supporters and opponents of CBA commonly maintain that it can, and indeed at least sometimes will, do so. I argue in this paper that it is far from obvious that they are correct.

Here, in short, is the central issue. CBA involves translating all gains and losses that would result from a policy into monetary units. The death of an innocent person is counted as a cost, but the standard understanding of CBA holds that in certain cases the benefits of the policy change – of a looser regulation, for example – will be large enough as clearly to exceed this cost (and any others that are involved) and that CBA will thus recommend the policy change. I argue for two points. The first is that what CBA has historically been understood to require in order to approve of a project is frequently misrepresented in contemporary literature. The notion of compensation has historically played an important role in explaining and justifying the policy decisions that CBA recommends, but this notion is often altogether absent from contemporary discussions. CBA’s demand is that it must be possible for those who gain through a policy

---

<sup>5</sup> U. S. Environmental Protection Agency, “Clear Skies Act: Basic Information,” [www.epa.gov/air/clearskies/](http://www.epa.gov/air/clearskies/). As I understand them, these numbers are relative to no air pollution-related regulations whatsoever. Many would argue, of course, that the Clean Skies Act saves significantly fewer lives than certain competing policies would.

<sup>6</sup> Michael Ong, James M. Lightwood, and Stanton A. Glantz, “Health and Economic Impacts of the Proposed Florida Smokefree for Health Initiative,” *Tobacco Control Policy Making* (2003), Center for Tobacco Control Research and Education.

<sup>7</sup> Jocelyn Kaiser, “How Much Are Human Lives and Health Worth?” *Science* 299 (2003), pp. 1836-7.

<sup>8</sup> I intend no judgment as to whether “lethal policies” are good or bad. Why the “reasonably quickly” qualifier is present, and how it should be interpreted, will be commented upon later in the paper.

change to compensation those who lose through the change, but this is frequently and inappropriately characterized as the demand simply that the policy's benefits exceed its costs. The second point is that once CBA's demand in respect to compensation is accurately understood, the case for thinking that CBA can approve of lethal policies is much more difficult to make than it otherwise appears to be. Offhand, it is difficult to see how even the potential for compensation exists in a situation where any of those who lose through a policy change lose so badly as to die. I review the standard economic account of how CBA's demand for the possibility of compensation can be satisfied in respect to lethal policies: what the policy is (purportedly) doing is simply imposing a risk on members of a population and so all that needs to be – and what frequently is – compensable is exposure to this risk. I take several tacks in explaining why this approach is unconvincing, including one that involves the proposal of a way of conceiving of costs that is, I believe, inherently more compelling than the standard economic conception and that entails that the costs imposed by a lethal policy are (almost always) of a non-compensable sort.

## **II. The Importance of the Notion of Compensation**

Crucial for properly assessing CBA is properly understanding it; and that requires, among other things, appreciating the relationship between CBA and several concepts bearing the name of the Italian economist Vilfredo Pareto. A distribution of goods/services is said to be Pareto efficient or Pareto optimal if it is such that in order to make one person better off through a redistribution another person would have to be made worse off. Such a distribution is optimal in the sense that any alternate distribution would be counted by at least someone as inferior to the original distribution; and it is efficient in the sense that there is no “free lunch” available through

an alternate distribution. One possible criterion for evaluating proposed policies holds that a policy should be endorsed if and only if it promotes Pareto efficiency, that is, if and only if it would make at least one person better off and would not thereby make anyone else worse off. This is known as the Pareto Criterion (PC).

Before one can be in a position to assess this criterion, the notions “better off” and “worse off” have to be clarified. Teachers and practitioners of CBA typically maintain that these notions are relative to each affected individual’s preferences: one is better off exactly when one gets what one prefers; one is worse off exactly when one gets what one does not prefer. Now, this view is not without its problems. It is, for example, sometimes noted – more often by philosophers than by economists – that in order for “being better off” and “getting what one prefers” to hang together tightly, the preferences at issue must be fully informed and consistent with one another.<sup>9</sup> I will not pursue this issue here because economists typically do not insist on these conditions and my present aim is (at least in part) to tease out certain implications of the *standard* economic approach. I will, however, return to the point later in the paper. In any case, how much better or worse off a change would make an individual must be expressed in some standard unit – money is convenient enough – and this is done by asking the individual to specify how much money she would be willing to pay (to get a positive change or avoid enduring a negative change) or willing to accept (to endure a negative change or avoid getting a positive change).

To be sure, a policy that satisfies PC has something going for it. Especially if the initial distribution was a fair one, there is a powerful argument according to which no one can claim to have been wronged through a change that satisfies this criterion. This is no trivial thing when it

---

<sup>9</sup> I have myself elsewhere dealt at length with the issue of ill-informed and inconsistent preferences as they relate to individual well-being. Citation omitted for the sake of blind review.

comes to publicly justifying a policy decision. Nonetheless, approving *only* of those policy changes that satisfy PC turns out to be quite restrictive in practice, since almost any large-scale public policy will make at least someone worse off. For example, consider the problem of the pollution of air, water, and soil in and around shipping ports. In my home state, California, many jobs and homes are located near ports in big cities, including Los Angeles, Long Beach, Oakland, and Stockton. According to the *San Francisco Chronicle*:

Cargo ships using heavy bunker fuel are partly responsible for making ports major sources of pollution, worsened by fleets of diesel-powered trucks that carry cargo in and out of facilities. Numerous studies have associated fine particulate matter, mainly from diesel, with a variety of respiratory and cardiovascular problems, ranging from aggravated asthma to irregular heartbeats, heart attacks and early death in people with heart or lung disease.

Some West Oakland residents are exposed to roughly five times more diesel particulates than residents in other parts of the city. The particulates are 90 times more concentrated than the state average, according to a 2003 report by the Pacific Institute, a nonprofit advocacy group.

“Goods movement,” says the California Air Resources Board, is a cornerstone of the state’s economy. But its own study concluded that for 2005, pollution from the trade is estimated to be responsible for *750 of California’s estimated 9,000 premature deaths associated with air pollution.*<sup>10</sup>

Options exists for reducing this pollution: there are cleaner-burning fuels, and work could be turned over to smaller, cleaner engines on the ship once the ship is within a certain number of

---

<sup>10</sup> George Rain, “Shipping’s Dirty Cargo,” *San Francisco Chronicle*, March 4, 2006 (emphasis added). For considerably more on the issue of port pollution, I recommend consulting the Bluewater Network, available online at: [http://www.bluewaternet.org/campaign\\_ss\\_ships.shtml](http://www.bluewaternet.org/campaign_ss_ships.shtml).

miles of the port. Not surprisingly, however, such alternatives cost money for someone, somewhere.<sup>11</sup> For this reason, a proposal to tighten regulations governing pollution in and around ports would not seem to satisfy PC. And obviously proposals to loosen such regulations would impose costs – additional health costs – on those working and living in and around the ports, and so would also apparently be rejected by PC.

However, before concluding that PC must therefore reject any change to the status quo, notice that if (1) those who were made better off through a policy change were to pay those who were initially made worse off through the policy so that those who were originally made worse off now considered themselves to be as well off as they were before the policy change, and (2) after this compensation took place none of those who originally gained through the policy was worse off, and at least one of them was better off, than was the case prior to the enacting of the policy, then PC would, in a manner we might call “indirect,” be satisfied after all. Such a policy and subsequent redistribution would actually promote Pareto efficiency.

As it turns out, it is not the case that CBA licenses only those policies that satisfy PC (directly or indirectly). Rather, it licenses those policies that have the *potential* to satisfy PC. More properly put, CBA holds that a policy should be approved when its benefits outweigh its costs<sup>12</sup>, with the gainers being *able* to compensate the losers such that the losers are no worse off

---

<sup>11</sup> Who exactly will *ultimately* incur these costs is a complicated issue, but this distributional question is not important from the perspective of the Pareto Criterion. All that matters is whether at least one person would be made worse off by the change, in this case to tighten emissions regulations. The notions of costs and benefits are handled in the same way that the notions of better off and worse off are. That is, costs just are the amount of money that affected individuals would be willing to pay in order to avoid, or will to accept in order to endure, a change that is not preferred; and benefits are the amount of money that affected individuals would be willing to pay to get, or willing to accept in order not to get, a preferred change.

<sup>12</sup> Some will add “by a greater margin than would be the case on any other policy” (or something of this sort) here. This would be added to account for the fact that there might be multiple policies each of whose aggregate gains outweigh its costs, but each to different extents. According to Layard and Glaister, this addition is not necessary because “it seems quite natural to refer to the ‘benefits of the next best alternative to A’ as the ‘costs of A’. For if A is done, those alternative benefits are lost.” Richard Layard and Stephen Glaister, *Cost-Benefit Analysis*, (Cambridge, UK: Cambridge University Press, 1994), p. 2. In any case, the notions of “cost” and “benefit” are to be understood in the same way that the notions “better off” and “worse off” are. That is, costs just are the amount of

than they were originally and the gainers are still better off than they were originally. The aforementioned compensation scheme must be possible, but the *possibility* is taken as sufficing to justify the policy; the compensation need not actually be paid. This criterion is known as the Kaldor-Hicks Criterion (KHC), and when it is satisfied the corresponding distribution is described as Kaldor-Hicks efficient.<sup>13</sup> There are various reasons why those who endorse KHC as the appropriate policy criterion say that actual compensation payments are not demanded. I will here mention the one I find most compelling, which is also, as it turns out, the most pragmatic in nature. When cost-benefit analysts calculate costs and benefits, they frequently do so by soliciting willingness-to-pay (WTP) or willingness-to-accept specifications not from every single person affected, but rather from a random sample of those people affected. Statistical practices make it possible to generalize total costs and benefits from this sample; and it is a good thing, since directly contacting each affected individual would typically be extraordinarily expensive and time consuming. One result of this, however, is that precisely how much each affected individual would need as compensation is simply not known.<sup>14</sup>

The philosopher who is just learning the principles of CBA could not be blamed for wondering exactly how KHC differs from what philosophers in the utilitarian tradition have been saying for centuries. It would be very surprising if there were no difference, since this would make entirely mysterious why Kaldor and Hicks have a maxim named after them! The difference is not that KHC is used to evaluate public policies while utilitarianism is used to evaluate only

---

money that affected individuals would be willing to pay in order to avoid, or will to accept in order to endure, a change that is not preferred; and benefits are the amount of money that affected individuals would be willing to pay to get, or willing to accept in order not to get, a preferred change.

<sup>13</sup> The criterion is also known as the Potential Pareto Criterion or the Potential Compensation Criterion.

<sup>14</sup> That compensation is not demanded by KHC does not mean that KHC somehow forbids compensation or that policies endorsed by CBA will as a matter of fact never involve actual compensation. For an interesting discussion of compensation for the nearby siting of environmental hazards (such as waste disposal facilities), including reflection of whether such compensation should be of a monetary or an in-kind sort, see Edmundo Claro, "Exchange Relationships and the Environment: The Acceptability of Compensation in the Siting of Waste Disposal Facilities," *Environmental Values* 16 (2007), pp. 187-208.

“private,” individual actions. Utilitarianism has no doubt be understood by at least some to be a policy criterion. Instead, in short, the difference is in the demand that compensation be possible. For utilitarians it is the mere fact of gains outweighing losses that justifies a policy change; whereas according to KHC the justificatory work is done by the fact that the losers could be compensated. Even if it is true that the potential exists for winners to compensate losers any and every time gains exceed losses, one might endorse policies whose gains exceed their losses without thinking the justification for this endorsement has anything to do with the possibility of compensation. One who did so would be proceeding on utilitarian, not Kaldor-Hicksian, grounds.

Problematically, however, if one listens to many of those who write about CBA, including many of its most prominent teachers, there is in fact no difference between KHC and the utilitarian theory of right action. Consider the following characterizations of KHC:

- Richard Layard and Stephen Glaister say that KHC maintains that a policy or project is justifiable “if [its] benefits exceed those of the next best alternative course of action, and not otherwise.”<sup>15</sup>
- A. Myrick Freeman claims that KHC “asks whether the aggregate of the gains to those made better off measured in money terms is greater than the money value of the losses of those made worse off. If the gains exceed the losses, the policy is accepted by this criterion.”<sup>16</sup>
- According to Michael Munger:

The Kaldor-Hicks criterion might be thought of as having two parts:

---

<sup>15</sup> Layard and Glaister, *Cost-Benefit Analysis*, *ibid.*

<sup>16</sup> A. Myrick Freeman, “The Economic Basis of the Economic View of the Environment,” Center for the Study of Values and Social Policy at the University of Colorado at Boulder, as reprinted in *The Environmental Ethics & Policy Book* (2<sup>nd</sup> Edition), Donald VanDeVeer and Christine Pierce (eds.), (Belmont, CA: Wadsworth, 1998), pp. 293-301 ( p. 296).

- Compare two states of the world,  $S_1$  and  $S_2$ . If one is Pareto superior, no further analysis is needed: the Pareto superior one is preferred.<sup>17</sup>
  - If both states are Pareto optima, add up the gains and losses to all citizens from choosing each alternative. Select the policy that maximizes the difference between the gains to the gainers and the losses to the losers. Importantly, it is not required that the gainers compensate the losers, but only that the excess value created by the policy choice is maximized.<sup>18</sup>
- Hank Jenkins-Smith says that “in the comparison of policies, the relevant criterion” – which he earlier has explicitly identified as KHC – “is: which policy option serves to create the *largest net gain in social well-being?*”<sup>19</sup>

More examples could be cited. What is important to recognize is that while the conceptual connection between PC and KHC is not contested – Munger acknowledges it explicitly, and even those who do not will admit that the latter is often referred to as the Potential Pareto Criterion – the aforementioned characterizations do not sufficiently reflect this connection. They do not appreciate the significance of, and do not appropriately reflect, KHC’s demand that *compensation* be possible. So, by stripping mention of compensation from their characterizations, these authors tacitly rob KHC of its distinctiveness and of much of what those

---

<sup>17</sup> Note: In comparing two distributions,  $A$  and  $B$ ,  $A$  is called Pareto superior to  $B$  if in  $A$  at least one person is better off, and no one is worse off, than in  $B$ .

<sup>18</sup> Michael C. Munger, *Analyzing Policy: Choices, Conflicts, and Practices*, (New York: W. W. Norton and Co, 2000) p. 103.

<sup>19</sup> Hank Jenkins-Smith, *Democratic Policies and Policy Analysis*, (Pacific Grove, CA: Brooks-Cole Publishers, 1990), p. 22 (emphasis in original). Actually, he first says: “The Kaldor-Hicks criterion allows redistributions that increase net welfare such that those who gain from the distribution could compensate those who lose, restoring the losers to their prior level of well-being, while the winners retain enough of their gains to be better off than they would have been without the redistribution.” However, he then follows this with the claim I have quoted in the body of the paper. My claim, which I hope to make convincing, is that the latter characterization is not identical to the former; it is not the case that the latter is satisfied anytime the former is.

who developed CBA's conceptual underpinnings saw as recommending KHC as a policy criterion.<sup>20</sup>

### III. Death and the Possibility of Compensation

An additional problem with the aforementioned characterizations of KHC is that they obscure just when it is – if ever – that CBA approves of lethal policies. Consider the following case, which is simplistic in order to illustrate a point. The state is considering a policy that will affect ten people. Through the policy, nine will immediately benefit, according to their own WTP, by a measure of \$1 million each.<sup>21</sup> The tenth, however, would certainly die immediately and without any warning whatsoever. The benefits of this policy are easy enough to figure: they are \$9 million. The costs are more difficult to figure, and the coming pages will contain much discussion of how this is and should be done. For now, let us simply note that economists who work in this area these days typically associate a cost of roughly \$6 million with each loss of human life—though (1) different economists may suggest slightly different numbers, (2) the number used will (in the United States at least) vary from and agency to agency, and (3) in theory the economic rationale for using that same single figure regardless of whose life it is (what that person is like, what her/his personal WTP would be, etc.) is weak.<sup>22</sup> So, let us for now call this the cost of this policy. We see, then, that the benefits of this policy would outweigh the

---

<sup>20</sup> Whether the potential for compensation really does have any special justificatory power is an issue that will be taken up later in the paper.

<sup>21</sup> It is of course unlikely that any actual policy will kill anyone immediately; but there should be nothing illicit about imagining a policy that would do so. This move simply helps us focus on the important issue.

<sup>22</sup> Note that I am careful not to say that this \$6 million represents the value of a human life. While it seems to me that economists are indeed trying to put a price on human life, many economists insist on different language: they say that they are merely valuing a “statistical” life. In any case, while different economists and agencies suggest different monetary amounts, I will use \$6 million of the sake of roundness in examples. The literature in this area is enormous. A good start might be with W. Kip Viscusi, *Rational Risk Policy: The 1996 Arne Ryde Memorial Lectures*, (Oxford: Clarendon Press, 1998), and W. Kip Viscusi, “The Value of Risks to Life and Health,” *Journal of Economic Literature* 31 (1993) pp. 1912-1946. See Alan Krupnick, “The Value of Reducing Risk of Death: A Policy Perspective,” *Journal of Policy Analysis and Management* 21 (2002), pp. 275-282.

costs. It would thus be approved of by KHC as the above authors characterize it. However, when we are prompted – as we will be by a proper characterization of KHC – to think about the potential for compensation, we note that it is hardly obvious that the person who is made to die through the policy can be compensated – even by that \$6 million – and in turn we see the claim that CBA will approve of this policy as more dubious. No doubt more rigor is in order, and I will indeed be more rigorous in coming pages; but the preceding example suffices to show that the aforementioned characterizations of KHC mask something important. Where KHC is described as demanding simply that benefits exceed costs, then, at least where economists are, as is typical, calculating costs simply by taking that (roughly) \$6 million figure and multiplying it by the number of lives that are expected to be lost by the policy change, one is apt to overestimate the ease with which CBA can approve of lethal policies.

It is interesting to note that this problem of compensation in cases where someone dies did not seem even to have occurred to Kaldor himself. In any case, he does not engage with it at all. In the seminal paper in which Kaldor laid the groundwork for the criterion which would eventually bear his name – and it seems to me that he laid much more of that groundwork than Hicks did – Kaldor is discussing something as relatively mundane as the repeal of the “Corn Laws.” These laws, in force for part of the 19<sup>th</sup> century in Britain, imposed tariffs on foreign grain import in order to protect the price of the domestic crop. Perhaps it is difficult to imagine anyone being made to die from changes to those Corn Laws—though to me not impossible. Even so, Kaldor fails to recognize, or at least explicitly to acknowledge, that not all cases are like this. He says:

In all cases, therefore, where a certain policy leads to an increase in physical productivity, and thus of aggregate real income, the economist’s case for the

policy is quite unaffected by the question of the comparability of individual satisfactions; since in all such cases it is *possible* to make everybody better off without making anybody worse off. There is no need for the economist to prove – as indeed he never could prove – that as a result of the adoption of a certain measure nobody in the community is going to suffer. In order to establish his case, it is quite sufficient for him to show that even if all those who suffer as a result are fully compensated for their loss, the rest of the community will still be better off than before.<sup>23</sup>

The port pollution case discussed above would seem to be one of many in which a policy change might lead to that “increase in physical productivity, and thus of aggregate real income,” but where, because the change results in death for some, it is not possible to make everybody better without making anybody worse off.

While Kaldor was not obviously concerned with compensation in connection with lethal policies, economists since then certainly have tried to deal with the problem. Plainly, the important question is how the cost figure is arrived at in connection with lethal policies. Economists maintain that even when it is known – as it frequently is – that a policy will lead to human deaths, it is almost never the case that it is known in advance exactly *who* will die. Rather, members of a population are subject to a certain risk of death. Thus, the question that economists hold to be the proper question to ask is (roughly, and in a generic form): “How much

---

<sup>23</sup> Nicholas Kaldor, “Welfare Propositions of Economics and Interpersonal Comparisons of Utility,” *The Economic Journal* 49 (1939), pp. 549-552 (p. 550); emphasis in original. Kaldor does recognize certain complications: in the first footnote (p. 551) he says, “An increase in the money value of the national income (given prices) is not... necessarily a sufficient indication of this condition being fulfilled: for individuals might, as a result of a certain political action, sustain losses of a non-pecuniary kind—e.g., if workers derive satisfaction from their particular kind of work, and are obligated to change their employment, something more than their previous level of money income will be necessary to secure their previous level of enjoyment; and the same applies in cases where individuals feel that the carrying out of the policy involves an interference with their individual freedom.” Nonetheless, he stops short of considering the problem I am pressing.

would you need to be paid in order to accept such-and-such an increase in your risk of dying?”<sup>24</sup>

This way of thinking suggests that we should understand a lethal policy to be one whose risks make it statistically likely that more people will die than are expected to die by another policy

(given the risks the various policies involve and the number of people exposed to those risks).

The claim will then be that all that needs to be compensable is facing whatever the relevant risk

of death. If this is right, then the question is: For each individual affected by the policy – and

more to the point, each individual on whom the policy imposes a risk of death – is there some

amount of money that the individual would accept in order willingly to face that risk? If so, and

if the policy is expected to generate benefits that exceed the sum of the amounts specified by the

individuals facing this risk, then each person is in theory compensable—and this includes those

persons (or that person) who happen(s) to be statistically unlucky and actually succumb to the

risk and die.

Let us go back to the example of the policy that kills one person and benefits nine. Tweak

the example slightly so that rather than it being the case that one person will certainly die

immediately and without warning, the policy would impose a cost on each of ten different people

by raising each of these people’s chance of immediate death from 0% to 10%. This increased

risk of death for these ten people is the only cost associated with the policy; and there is no

benefit for them associated with the policy. Imagine that each of these ten people is willing to

---

<sup>24</sup> The question, “How much would you be willing to pay to avoid getting such-and-such an increase in your risk of dying?” is supposed to be essentially identical. In fact, it is not. It is well known that the amount people are willing to accept to endure risks is higher than the amount they are willing to pay to avoid the risk. In part, this is because one’s willingness to pay is constrained by ability to pay in a way that one’s willingness to accept is not. It is also (at least suspected to be) because the willingness to accept question suggests that one has a right not to be exposed to the risk whereas the willingness to pay suggests otherwise. One other difference between the willingness to accept and willingness to pay formulations – and one that is not well appreciated, I think – is that the amount one indicates one is willing to accept in order to face the risk is easier to construe as an amount that would compensate one than is the amount one indicates one is willing to pay to avoid facing the risk. If I say I am willing to pay \$X to avoid such-and-such a risk and the risk is imposed on me, it hardly seems that if I am not made to pay \$X I therefore count as having been compensated for facing the risk.

accept this risk for \$600,000. The total cost of the policy is thus \$6 million. Economists have two general strategies for ascertaining this kind of information. Sometimes the amount people would be willing to accept to face (or pay to avoid) the risk is gleaned by looking at their actual behavior—for example, how much more individuals who work in a risky job are paid than those who work in a safer (but in other respects similar) job. This is known as the revealed-preference method for ascertaining individual preferences. Other times, the amount affected individuals are willing to pay is found through the use of survey devices. This is known as the stated-preference method or contingent valuation. However the \$600,000 figure was arrived at vis-à-vis each of the ten people, let us suppose that the policy would deliver \$9 million of gross benefits to an 11<sup>th</sup> individual; and it would impose no costs on this person.

In this case, the benefits would seem to outweigh the costs: again, \$9 million in benefits versus \$6 million in costs. If Layard and Glaister, Freeman, Munger, and Jenkins-Smith are correct, then KHC is satisfied. But is it really? As I have said, it actually depends on whether the losers are compensable. In the case we have imagined, the 11<sup>th</sup> could seemingly give each of the other ten \$600,000 and still be better off than she (the 11<sup>th</sup>) was originally. The ten would each then have received what each said she would require in order to accept exposure to the higher risk of death. But of course statistically the chances are that the policy will actually result in one of them dying—and, again, in order that we avoid certain complications, we are imagining that this would happen immediately upon enactment of the policy change. Let us thus suppose that this happens: one of the 10 who were at risk – let us call her Ann – succumbs to the risk and dies. The question, then, is: Does the \$600,000 count as sufficient to compensate Ann for her death? It may seem that if Ann claimed (as per the example) that she would accept this increased risk of death for \$600,000, then \$600,000 *must* count as sufficient compensation if Ann succumbs to

that risk and dies. Imagine that someone says: “I will spend \$5 on a lottery ticket for an infinitesimally small chance to win \$10,000,000, but only if I do not lose!” This would be ridiculous – the ticket is only worth a small sum (like \$5) precisely because it has such a small chance of winning – but this is essentially what Ann’s claim would be if the \$600,000 did not count as compensating her. A risk is no risk at all, and so not worthy of any compensation whatsoever, if succumbing to that risk is not a genuine possibility. Thus, in spite of Ann’s death, the policy change would seem to satisfy KHC.

This proffered explanation of how it is that CBA can at least sometimes approve of lethal policies is a thoughtful one. But it is by no means clearly correct. Admittedly, it was not known in advance who would die, but this does not entail that everyone must be asked “How much would you need to be paid in order to accept such-and-such an increase in your risk of dying?” So long as we are confident that someone will die, there is no reason to insist that asking each person “How much would you be willing to accept in order to die?” is inappropriate. Naturally, the latter question and the former one are likely to generate different responses. Now, I would not say that it is impossible that there could be someone who *could* straightforwardly be compensated even for her own *immediate* death. Imagine a very loving but extremely desperate parent who agrees to be killed in exchange for a large sum of money that will benefit her sick and starving children.<sup>25</sup> If the agreed upon amount of money is delivered, this person may indeed count as compensated for her death. It seems clear, however, that such cases are going to be extremely odd and rare. Most people, if asked how much they would need to be paid in order voluntarily to accept being killed would probably think the question itself was crazy and unintelligible. There is no amount of money the average person would specify as compensating her for her death, especially if the death took place very soon, not many years down the line.

---

<sup>25</sup> Perhaps her organs could uniquely benefit a very wealthy but extremely ill person.

The qualification in the previous sentence (which is a qualification I have inserted at various earlier points in the paper) is an important one. To wit, it is not hard to believe that if one knows, or at least expects, that the policy-induced death will not come for many, many years, there is an amount of money that if paid out now would compensate this person – by her own account, even – for her nonetheless-earlier-than-it-would-have-been death. For example, if I believe that a policy will create conditions whereby I will die 2 years earlier than I would have – at age 95 rather than 97, say – a payment of \$6 million now (when I am still many years younger than 95) would probably be one I would count as sufficient to compensate me. This is why my definition of “lethal policy” was initially specified as it was – namely as a policy that would *reasonably quickly* lead either to deaths when another available policy would not, or to more deaths than would another available policy – rather than simply as any policy that would lead to more deaths than an alternate policy would. That is, when I talk about lethal policies, I have in mind those in which the death would come about fairly soon—soon enough, let us say, where no one other than the likes of the extremely desperate parent mentioned above would consent to the policy for a finite amount of money.

Let us return to Ann. While it is true that Ann indicated that she would face a particular increased risk of death for \$600,000, to say that in the end all that happened is that Ann faced this risk is to say something false. That she “succumbed to this risk” is true, but also euphemistic. We must be frank: Because of the conditions the policy change created, Ann *died*. The loss imposed on Ann must include the loss of her life. It is not just for facing a risk, but rather for the loss of her very life, for which she should be understood as needing to be compensated; and her life is worth more to her than \$600,000. The \$600,000, therefore, does not compensate her for the change the policy created. In fact, if, as it is reasonable to suppose, Ann would not have

consented to dying immediately for \$6 million, then even the full \$6 million would not compensate her for her immediate death. Her life is worth more to her than even this. Regardless of what amount she specifies herself as needing in order willfully to face the risk, it is hard to see how *any* amount of money would be sufficient to compensate her when she dies immediately upon (or even reasonably soon after) enactment of the policy change. And it is in turn hard to see how such a policy could satisfy KHC.

At this point it may appear that we have two reasonable descriptions of Ann's situation that generate contradictory conclusions, which would mean that in connection with lethal policies KHC is plagued by a paradox. It has been claimed that this has everything to do with the distinction between *ex ante* and *ex post* perspectives, and more specifically that the suspicion that someone like Ann would not be compensated by the \$600,000 arises only when one takes the latter perspective.<sup>26</sup> What exactly is the difference between these? Obviously, *ex ante* evaluation is before the fact, and *ex post* is after the fact. But there is an ambiguity still. The *ex post* perspective is sometimes understood to be available only after the policy change has actually occurred and analysts are looking back to determine if the change was, after all, justified. However, if it is understood in this way, this perspective is a fairly silly one to take. After all, what good does it do to evaluate a policy decision after it the relevant change has taken place? If an evaluation is going to be done at this point, it is best conceived of as an evaluation looking *forward* in order to determine whether some additional policy change (which might involve a change back to the original condition) is justified. So, the *ex post* perspective is frequently instead understood as one taken by people evaluating a policy change *before* it has gone into effect, but where they imagine what the preferences of affected individuals will be *after* the

---

<sup>26</sup> See especially J.M. Buchanan and R.L. Faith, "Trying Again to Value a Life," *Journal of Public Economics* 12 (1979), pp. 245-248; as well as Alastair Ulph, "The Role of Ex Ante and Ex Post Decisions in the Valuation of Life," *Journal of Public Economics* 18 (1982), pp. 256-276.

change occurs. The distinction is also sometimes understood such that knowledge of who precisely will die is available from the *ex post* perspective and that knowledge of who faces what risks is available from the *ex ante* perspective.

However precisely these perspectives are understood, I do not believe that much is resolved by talking about them. For one, even if it were true that the paradox arises out of the existence of these two perspectives, noting this would not help resolve the paradox if – as so far seems to be the case – both perspectives seem legitimate. Now, M.W. Jones-Lee would have us believe that the *ex ante* perspective is appropriate because it is the one recommended by the “value judgment that social decisions should reflect private tastes and preferences... as these exist at the time of the social decision.”<sup>27</sup> Obviously, this is contentious. However, even if the “value judgment” Jones-Lee appeals to is compelling, that is, even if the *ex ante* perspective is the appropriate one, in no way does this show that it is mere exposure to lethal risk, not death itself, that must be compensable. This is because those tastes and preferences that exist even *at the time of the social decision* surely include, for each person who faces a lethal risk through the project, a preference not to die! Nor do we – to return to our earlier example – need to know ahead of time that Ann is the one among those facing the risk who will actually die in order to entertain the suspicion that the policy change is of a non-compensable sort. All we need in advance to be confident of, and what are in advance confident of in the example involving Ann and with respect to many real-life policies that bear on human mortality, is that at least *someone* will die through the policy. When this is true, we can reasonably ask whether that person who will die, whoever she or he is, is compensable for what will happen. And the answer, in essentially every case where that death comes reasonably soon, will be “no.”

---

<sup>27</sup> M.W. Jones-Lee, “Trying to Value a Life: Why Broome Does Not Sweep Clean,” *Journal of Public Economics* 12 (1979), pp. 249-256.

For these reasons, I do not think the apparent paradox truly is explained by the *ex ante/ex post* distinction. In the end, in fact, I think the appearance of a paradox is just that—an appearance. We do not need to call the description according to which Ann’s loss consists merely in exposure to risk and so is compensable a “reasonable” one. The discussion of the *ex ante* and *ex post* perspectives is really getting at the issue of how we should understand costs. This is a very important issue, but it seems to me that we can fruitfully bypass the distinction between these perspectives altogether by settling on a plausible account of well-being and letting *this* shape our understanding of costs (and benefits, too). Earlier I mentioned that a person’s being made better off and that person’s preferences being satisfied do not always hang together if she is not well informed when she forms and specifies her preferences. Consider, for example, someone who pulls up to a toll bridge willing to pay \$5 to cross. Imagine, though, that the bridge – perhaps unbeknownst to everyone – is going to collapse momentarily. It is utterly implausible to say that crossing will make this person better off by a measure of \$5. No, it will not—it will make this person dead! If one is not well informed, there is no guarantee that getting what one wants makes one better off. One may say at this point that all I have done is shift the *ex ante/ex post* question over to the concept of well-being—that is, that whether this person is better off by a measure of \$5 depends, again, precisely on whether one takes the *ex ante* or *ex post* distinction. This is not true. What it depends on is whether or not one goes for an account of well-being that is purely subjectivist. And it seems clear that a *purely* subjectivist account of well-being, such as the one standardly appealed to in economics, is deeply flawed.

Here is another example, which will lead us to (a rough sketch of) that more plausible, quasi-subjectivist account of well-being and in turn to the appropriate conception of costs. Imagine someone who has very high cholesterol but prefers to eat fast food frequently.

Intuitively, it seems that getting what she prefers will make this person worse off, and so the existence of such a person poses a challenge to the attempt to link getting what one prefers with being made better off. However, if she does not know that she has high cholesterol, that high cholesterol threatens her longevity, or that frequently eating fast food will further elevate her cholesterol numbers, it suddenly seems even easier to deny that satisfying the preference for fast food does not enhance her well-being. On the other hand, it is more difficult to deny this when she knows these various facts and still prefers the fast food. And it would be still more difficult to deny if, in addition to these facts, she knew that eating the fast food would lead to her suffering a fatal heart attack sometime soon and nonetheless still preferred the fast food. Certainly such a preference is difficult to make sense of, but if she genuinely had it, I would be strongly tempted to admit that the fast food contributes to her well-being. Here we have to say that well-being does not always consist in better health, *ceteris paribus*, but that may well be the most plausible view—especially given the problems with so-called objective-list and teleological accounts (which will not be discussed here). This is at least a start to a plausible account of well-being: it is not purely subjectivist, but in that it nonetheless makes an important role for preferences, it can be called quasi-subjectivist. Well-being, on this account (which may be called an informed-preference account) consists in what one *would* prefer *if* one were genuinely well informed.

The proper conception of costs falls directly out of this account. That is, it follows from the fact that when one is ill informed, one is not necessarily better off when one gets what one prefers, that we should understand costs as what affected individuals would be willing to accept in order to endure the negative change if they had complete information. Complete information consists in accurately understanding the risk one faces, but also in knowing whether one will

succumb to that risk. If each person on whom a policy change imposes an increased risk of death knows in advance whether or not she would succumb to that risk (when, etc.), then there is little doubt that the amount each person specifies herself as needing in the way of compensation for exposure to this risk *would* genuinely compensate her when she succumbs to the risk.<sup>28</sup> It seems reasonable to suppose that nearly any fully-informed individual who will lose her or his life through a policy change would – at least where this takes place soon, not decades into the future – not endure the policy change for any finite amount of money. In turn, the costs of the policy change will not be finite and no finite amount of benefits will suffice to satisfy KHC, which means that lethal policies will not be approved of by CBA.

Is this, combined with the earlier ruminations, enough to demonstrate that CBA must reject any lethal policy, or is such a conclusion too radical? The view that the only acceptable policy that prevents as many deaths as is technologically possible is unquestionably an extraordinarily demanding one: it means, for example, that if technology exists for preventing one more death, but this will cost – to be dramatic – \$1 billion, the technology must be used and the death must be prevented. Such a demand is, it seems to me, so demanding as to be indefensible, though I will not argue for this claim here. Perhaps KHC, even when properly interpreted, does not make quite this demand, since the unavailability of that \$1 billion for other purposes may entail that some altogether different kind of project that would have entailed the loss of fewer lives than its competitors is now not able to be funded. In any event, it is a clear implication of my argument that KHC is very a demanding criterion. Lethal policies should not, by the lights of KHC, be endorsed as frequently as they are now. If proponents of the use of

---

<sup>28</sup> Of course, if each individual knows what will happen to her when she specifies her WTP, she cannot properly be said to be facing a “risk”; but this is mere semantics. The “risk” language can be avoided if it is problematic.

CBA as an environmental policy decision mechanism believe that lethal policies are frequently the right policies, they should demand that CBA employ a new decision criterion.

That criterion might be one that is more genuinely utilitarian than KHC is. That is, it might be one that says simply that it is permissible to inflict extreme suffering or death on some small number of people if doing so brings about sufficient benefit for others. The challenge then, however, will be to justify this criterion; and the justification will not be able to refer to compensation or compensability in the way that KHC does. This, I think, is not a big setback, since the role that KHC gives to the notion of compensation is frankly not a terribly morally compelling one. After all, if I am one who loses through a policy change, why should I accept my losses as justified by the fact that the winners *could compensate me* and still be better off than they were originally? I believe I can understand why straightforward preference utilitarians say that all that matters is that benefits outweigh costs (making thereby no mention at all of compensation); though utilitarianism's indifference to distributional considerations surely counts as a serious drawback. I think I can also understand the demand that compensation actually be paid out from winners to losers; though I am sympathetic to the fact that cost-benefit analysts' tendency, in practice, to solicit WTP amounts from random samples of affected individuals interferes with their ability to identify precisely how much each individual would need to be compensated and so makes the demand for actual compensation difficult to satisfy in practice. I am also uneasy with the idea that if, say, the government tightens regulations on the shipping industry so that cleaner-burning fuel must be used, workers in and residents living near a port would be required to compensate the shipping industry for, in effect, not poisoning them! Nonetheless, my recognition of legitimate concerns on each of the two sides does not suffice to

recommend KHC, the awkward hybrid that says that a policy change is justified precisely by the potential for compensation.

In the end, though, nothing much is at stake for me when I suggest that the mere potential for compensation does not carry any justificatory weight. Even if that is mistaken, I believe this paper has done something important by showing, first, that modern commentators are unfaithful to KHC's real demand by reducing it to the requirement simply that costs outweigh benefits, and, second, that KHC, properly interpreted, and therefore CBA insofar as it employs KHC, cannot endorse many of the lethal policies that CBA's proponents have viewed as so obviously justified according to CBA's tenets.