

ARTICLE

# Not All Killings Are Equally Wrong

Todd Karhu

London School of Economics and Political Science  
[t.karhu@lse.ac.uk](mailto:t.karhu@lse.ac.uk)

(First published online 24 May 2019)

## Abstract

Many people believe that the wrongness of killing a person does not depend on factors like her age, condition, or how much she has to lose by dying – a view Jeff McMahan calls the ‘Equal Wrongness Thesis’. This article argues that we should reject the Equal Wrongness Thesis on the basis of the moral equivalence between killing a person and knocking her unconscious.

## I. Introduction

Ronald Dworkin once observed that ‘most people think ... that the murder of a depressive handicapped octogenarian misanthrope is as heinous, and must be punished as seriously, as the murder of anyone younger or healthier or more valuable to others’.<sup>1</sup> Jeff McMahan, a prominent defender, calls this view the *Equal Wrongness Thesis*.<sup>2</sup> As he puts it, the moral objection to killing persons ‘does not vary with such factors as the degree of harm caused to the victim, the age, intelligence, temperament, or social circumstances of the victim, whether the victim is well liked or generally despised, and so on’.<sup>3</sup>

Whether the Equal Wrongness Thesis is true is of obvious theoretical interest. It also has important practical implications, since any factor (age, for example) that is irrelevant to the wrongness of killing is *ipso facto* irrelevant to its overall permissibility. Suppose we must decide, as the German Federal Constitutional Court once did, whether to pass legislation permitting the gunning down of hijacked passenger planes when doing so is necessary to prevent their use as a terrorist weapon to kill many more people.<sup>4</sup> If how wrong it is to kill a person is unaffected by how bad it is for her to die, then it makes no difference to the permissibility of shooting down these aircraft that the passengers on board will almost certainly die soon anyway.

<sup>1</sup>Ronald Dworkin, *Life's Dominion: An Argument about Abortion, Euthanasia and Individual Freedom* (New York, 1993), p. 85.

<sup>2</sup>Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life* (Oxford, 2002). See also Jeff McMahan, ‘Killing and Equality’, *Utilitas* 7 (1995), pp. 1–29.

<sup>3</sup>McMahan, *The Ethics of Killing*, p. 235.

<sup>4</sup>Bundesverfassungsgericht, *Judgment of the First Senate of 15 February 2006*. Of particular relevance is paragraph 132, in which the court argues that ‘[h]uman life and human dignity enjoy the same constitutional protection regardless of the duration of the physical existence of the individual human being’.

Kasper Lippert-Rasmussen has argued that we should reject the Equal Wrongness Thesis on the basis of the moral equivalence of killing someone and knocking her unconscious. Several authors have raised objections to his argument for their equivalence, however, and, as I shall show, there are other issues with his approach that have not yet been noted. Nevertheless, in what follows I shall argue that killing someone and knocking her unconscious are indeed morally equivalent in the way that is required to show that not all killings are equally wrong.

The structure of this article is as follows. In section II, I clarify the Equal Wrongness Thesis. In section III, I introduce Lippert-Rasmussen's argument against that thesis and discuss some of the ways I believe it falls short. In sections IV and V, I defend the claim that, other things being equal, it is no less wrong to deprive someone of a certain amount of conscious life by knocking her unconscious than it is to do so by killing her, regardless of when in her lifetime it takes place. I conclude, in section VI, by explaining why that premise is sufficient to defeat the Equal Wrongness Thesis. In brief, along with the Equal Wrongness Thesis, it implies that knocking someone temporarily unconscious for a short period is just as wrong as killing someone with many good years left to live. Because that result is absurd, we should reject the Equal Wrongness Thesis.

## II. Preliminary remarks

Two initial clarifications are necessary. First, as McMahan explains, the sense of 'wrongness' that is relevant for the Equal Wrongness Thesis is one that admits of degrees.<sup>5</sup> An act's wrongness in this sense is the degree to which it is morally objectionable: one act is more wrong than another in so far as it is subject to a moral objection that is harder to outweigh by countervailing moral considerations. Note that the Equal Wrongness Thesis applies only to acts of killing considered on their own. Suppose that on Monday I kill a young person and on Tuesday I kill an old person. The Equal Wrongness Thesis tells us that what I do on each day is equally wrong and would have been equally difficult to justify. But two acts being equally hard to justify does not imply that in a choice between them we ought to be indifferent. Suppose we must choose between killing a young person and killing an old person, for example as a necessary side effect of averting the deaths of hundreds of other innocent people. It is consistent with the Equal Wrongness Thesis that in this case we ought to kill the old person.<sup>6</sup>

Second, following McMahan and others, I understand the Equal Wrongness Thesis to be restricted in scope to a certain class of killings.<sup>7</sup> It does not apply to killings that differ in morally relevant but extrinsic ways, such as their effects (positive or negative) on third parties, or the relationship between the perpetrator and the victim. Nor does it apply to killings that differ with respect to the presence or absence of potential justifying

<sup>5</sup>McMahan, *The Ethics of Killing*, p. 190. See also Kasper Lippert-Rasmussen, 'Why Killing Some People is More Seriously Wrong than Killing Others', *Ethics* 117 (2007), pp. 716–38, at 717. Cf. Carlos Soto, 'Killing, Wrongness, and Equality', *Philosophical Studies* 164 (2013), pp. 543–59, at 551–4.

<sup>6</sup>McMahan, *The Ethics of Killing*, p. 237.

<sup>7</sup>McMahan, *The Ethics of Killing*, pp. 236–7. Other discussants of the Equal Wrongness Thesis also apply these restrictions to it. See Lippert-Rasmussen, 'Why Killing Some People', pp. 718–19; Soto, 'Killing, Wrongness, and Equality', p. 544; Matthew Hanser, 'The Wrongness of Killing and the Badness of Death', *The Oxford Handbook of Philosophy of Death*, ed. Ben Bradley, Fred Feldman and Jens Johansson (Oxford, 2013), pp. 392–409.

considerations, such as the victim's moral responsibility or her free consent to being killed. It does not apply to killings that differ in their modes of agency – for instance, whether the victim is killed as a means to some end or as a side effect of the pursuit of that end. Lastly, it concerns only the killing of individuals who satisfy the conditions of personhood, whatever those conditions are. The thesis is compatible with any view about the wrongness of killing non-persons. I will usually omit the foregoing qualifications and take them to be understood when I discuss instances of killing.

### III. Lippert-Rasmussen's trilemma

In 'Why Killing Some People Is More Seriously Wrong than Killing Others', Kasper Lippert-Rasmussen challenges the Equal Wrongness Thesis by way of a trilemma consisting of the following three claims:<sup>8</sup>

- (1) *The Equal Wrongness Thesis*: The degree to which different killings of persons are wrong does not vary: all such killings are equally wrong.
- (2) *The Unequal Wrongness of Renderings Unconscious Thesis*: For any period  $\omega$ , it is more wrong to render a person unconscious for a period longer than  $\omega$  than it is to render a person unconscious for  $\omega$ , other things being the same.
- (3) *The Equivalence Thesis*: It is neither more, nor less, wrong to deprive a person of a certain amount of conscious experience by killing her than it is to deprive her of the same amount of conscious experience by rendering her unconscious, other things being the same.

To see that these three claims cannot all be true, consider the following two pairs of acts. In the first pair, I kill someone with fifty years left to live and then kill someone with just one year left to live. According to (1), these two acts are equally wrong. In the second pair, everything is the same but instead of killing them I administer to each a powerful drug that induces a total coma that will last until her death. According to (3), what I do to the person who will live another fifty years is more seriously wrong than what I do to the one who will live for just another year. Yet if the two foregoing judgements are correct, then (2) cannot be true. For if depriving someone of a certain amount of conscious experience by killing her is no more or less wrong than doing so by knocking her unconscious until her death, the relative wrongness must be constant across both pairs. Lippert-Rasmussen argues that of these three claims we should give up the Equal Wrongness Thesis.

I take the crucial premise in Lippert-Rasmussen's argument to be (3). There are, however, a number of ways to challenge this premise. First, unlike when she is knocked unconscious, when a person is killed she will never regain consciousness. Since it may matter for the wrongness of an act whether it deprives its victim of her last period of conscious life, we might doubt that depriving someone of a given amount of conscious experience by knocking her unconscious is morally equivalent to killing her.

Lippert-Rasmussen suggests two individually sufficient ways of responding to this worry. One is to restrict (3) so that it covers only instances of unconsciousness that persist until the victim's death.<sup>9</sup> (Whether we take this restriction to be covered by a suitable interpretation of the premise's *ceteris paribus* clause or else to constitute a revised

<sup>8</sup>Lippert-Rasmussen, 'Why Killing Some People'.

<sup>9</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 727.

version of the premise itself is not a substantive matter.) Then it cannot be argued that the two types of act are morally different on the grounds that it makes a difference whether or not the victim will regain consciousness.

The other response is to maintain a more expansive, unrestricted version of (3) and argue that it is not in itself morally significant whether an act deprives its victim of what would have been her last period of consciousness. Lippert-Rasmussen proposes that it only seems more wrongful because we implicitly assume that depriving someone of what would have been her last conscious period is an especially great loss for her or is especially disrespectful.<sup>10</sup> With respect to the former, he notes that there is no reason, in principle, why the final period of a person's life is more valuable or important than any other. Yet this overlooks the possibility, argued for at length by Francis Kamm, that losing consciousness forever is non-comparatively bad.<sup>11</sup> It is also possible that losing what would have been one's last period of conscious existence, though it does not make the victim's life worse overall, is nonetheless worse for her relative to the time it occurs, because the loss of a given amount of future good matters more the less good one has in one's future. Moreover, it is not far-fetched to think that ensuring that someone will never again experience anything is a particularly egregious display of disrespect. Daniel Cohen and Morgan Luck have argued that the wrongness of a deprivation of consciousness depends on the proportion of the victim's remaining conscious life that it removes.<sup>12</sup> Before we are justified in accepting the more expansive version of (3), we need convincing replies to these objections. Ideally, we would also have a strong, positive argument for the view that the wrongness of knocking someone unconscious does not depend on when in the victim's life it occurs.

Because of these issues with the more expansive construal of (3), perhaps the opponent of the Equal Wrongness Thesis should restrict (3) to cover only those knockings unconscious that last for the rest of the victim's life. Lippert-Rasmussen argues that restricting (3) in this way is sufficient to answer the worry at hand, since the restrictive version of the trilemma that emerges is just as forceful as the unrestricted one.<sup>13</sup> But even if that were correct, it may be doubtful that depriving someone of consciousness by killing her is morally equivalent to doing so by knocking her unconscious until her death. Lippert-Rasmussen offers several arguments for their equivalence, but, as I will argue in section IV, none is entirely convincing. There are also disanalogies between killing someone and knocking her unconscious that he does not consider. For example, Matthew Hanser has since objected that the fact that killing someone terminates all her basic biological operations is a factor in making killing wrong.<sup>14</sup> Alternatively, it might matter simply that by killing someone, but not by knocking her unconscious, one causes the victim to cease to exist.

There is, however, a deeper problem with restricting (3) to cover only deprivations of consciousness that continue to the end of the victim's life. If we take this route, then in order for the trilemma to go through, (2) must also be understood to apply only

<sup>10</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 726.

<sup>11</sup>F. M. Kamm, *Morality, Mortality: Volume I: Death and Whom to Save From It* (Oxford, 1993). See also F. M. Kamm, 'The Purpose of My Death: Death, Dying, and Meaning', *Ethics* 127 (2017), pp. 733–61. Shelly Kagan also defends the view that death is non-comparatively bad in his *Death* (New Haven, 2012), pp. 278–9.

<sup>12</sup>Daniel Cohen and Morgan Luck, 'Why A Victim's Age Is Irrelevant When Assessing the Wrongness of Killing', *Journal of Applied Philosophy* 26 (2009), pp. 396–401.

<sup>13</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 727.

<sup>14</sup>Hanser, 'The Wrongness of Killing'.

between knockings unconscious that last for the rest of the victim's life. But the original, unrestricted version of (2) may have been compelling in the first place because of an implicit assumption that the instances of making people unconscious do not involve ending the victim's conscious life forever. Lippert-Rasmussen states that he will not pursue an independent argument for (2).<sup>15</sup> But for those who, like me, believe the Equal Wrongness Thesis has considerable *prima facie* plausibility, it may not be obvious that the wrongness of rendering someone permanently unconscious depends on the amount of conscious life she loses. Since the force of the trilemma lies precisely in the independent plausibility of the Equal Wrongness Thesis, if we restrict (3) to cover only acts that make their victim permanently unconscious there is probably insufficient reason to resolve the trilemma by rejecting the Equal Wrongness Thesis, rather than by rejecting (2).

To defeat the Equal Wrongness Thesis, therefore, we need to show both that the wrongness of knocking someone unconscious does not depend on when in the victim's life the unconsciousness occurs *and* that any other properties of killing someone that distinguish it from knocking unconscious – such as the fact that doing so ends the victim's life – are not themselves morally relevant.

The rest of this article is an attempt to do precisely that. In particular, I shall argue for the following claim:

- (C) It is as wrong to deprive someone of a certain amount of conscious life by knocking her temporarily unconscious in the middle of her life as it is to deprive her of the same amount of conscious life by killing her, other things being the same.

My strategy is first to defend, in a novel way, the restricted version of (3) – namely, that killing someone and knocking her unconscious until her death are morally equivalent, other things being the same. With that conclusion established, I will then offer a number of arguments to show that the wrongness of knocking someone unconscious for a given period is unaffected by when in the victim's lifetime it takes place. Together, these two premises entail (C), which is sufficient to refute the Equal Wrongness Thesis. That thesis implies that killing someone a month before she would otherwise have died is as wrong as killing someone with fifty good years left to live. Hence, if we accept (C), we can maintain the Equal Wrongness Thesis only by accepting that it is just as wrong to knock someone temporarily unconscious for a month as it is to kill someone with fifty good years left to live. Since that result is absurd, we should abandon the Equal Wrongness Thesis.

#### IV. The moral equivalence of killing and knocking unconscious until death

In this section I will defend my argument's first premise, which is the restricted version of Lippert-Rasmussen's Equivalence Thesis mentioned earlier:

- (A) It is as wrong to deprive someone of a certain amount of conscious life by knocking her unconscious until her death as it is to deprive someone of the same amount of conscious life by killing her, other things being the same.

<sup>15</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 722.

In his own defence of that thesis, Lippert-Rasmussen employs a pair of cases in which

- (1) I kill a person twenty years before he would otherwise have naturally died.
- (2) I knock a relevantly similar person unconscious for twenty years, at the end of which he dies a natural death, and I do so knowing that he will not regain consciousness before dying.<sup>16</sup>

Lippert-Rasmussen proposes, first, that these two acts are equally bad for their victims.<sup>17</sup> But the fact that two acts are equally bad for their victims does not imply that they are equally wrong. (Indeed, if an act's wrongness always correlated with its harmfulness we could reject the Equal Wrongness Thesis out of hand.) Lippert-Rasmussen acknowledges this, but argues that in the absence of an explanation for why two acts differ morally the fact that they are equally harmful is a reason to believe them equally wrong.<sup>18</sup> But if two acts seem to differ morally, then even lacking an explanation for *why* they differ, we are probably unwarranted in concluding them to be equally wrongful on the sole basis that they are equally harmful – and it is questionable that most people would intuitively believe that (1) and (2) are equally wrong. Moreover, there *are* a number of reasonable candidate explanations for that difference: for example, that killing someone ends that person's life, whereas knocking her unconscious does not.

The other argument Lippert-Rasmussen gives for the moral equivalence of killing someone and knocking her unconscious follows from existing theoretical accounts of *why* killing is wrong.<sup>19</sup> If the factors that explain why killing is wrong obtain in equal measure to analogous cases of knocking someone unconscious, then it seems reasonable to assume that the two types of act are morally equivalent. For example, if the intrinsic wrongness of killing is explained by its harmfulness, then, on the assumption that being knocked unconscious is just as bad for a person as being killed, it is reasonable to conclude that the two types of act are equally wrongful. Alternatively, suppose the wrongness of killing is fully explained by its disrespectfulness of the victim. Then, on the assumption that the factors that make killing someone disrespectful apply equally to knocking her unconscious, quite probably the two types of act are morally equivalent.<sup>20</sup> But whether or not all killings are equally wrong is itself a major consideration that any account of the wrongness of killing must accommodate, so appealing to theoretical accounts of the wrongness of killing to argue why it does not differ morally from knocking someone unconscious risks begging the question. Indeed, it is precisely

<sup>16</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 722.

<sup>17</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 722. He adds that any cost it would be permissible to impose upon the agent to prevent him doing (1) could also be permissibly imposed upon him to prevent him doing (2). I agree with Matthew Hanser ('The Wrongness of Killing', p. 379 n. 17) that this argument would sway only someone who already believes these acts to be morally equivalent.

<sup>18</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 722.

<sup>19</sup>Lippert-Rasmussen, 'Why Killing Some People', pp. 723–4.

<sup>20</sup>Lippert-Rasmussen, 'Why Killing Some People', p. 724. In defence of his claim that knocking someone unconscious is *ceteris paribus* no less disrespectful than killing her, Lippert-Rasmussen points to the fact that persons deserve respect not simply in virtue of being alive, but rather in virtue of 'certain cognitive and emotional capacities' ('Why Killing Some People', p. 724). But an unconscious person might retain her capacities for cognition and emotion. Furthermore, this argument seems to conflate the *bases* of respect for persons and the *manner* in which we should respect them. That we should respect people in virtue of certain capacities they have does not imply that respect for them consists in respect for those capacities.

on the basis of his conviction that all killings are equally wrong that McMahan rejects theoretical accounts of the wrongness of killing that imply otherwise, and it is also on the basis of that conviction that he proposes his own 'intrinsic worth' account, whose primary selling point is that it entails the Equal Wrongness Thesis.<sup>21</sup>

In the rest of this section I will attempt to establish the moral equivalence of killing someone and knocking her unconscious until death in a different way. My argument runs as follows. If, other things being equal, killing someone really is more wrong than knocking her unconscious until death, then killing must have some wrong-making property not shared by knocking unconscious until death. But there is no such property. Therefore, the two acts are morally equivalent.

I take it that the conceptual difference between killing someone and knocking her unconscious until death is that the former causes the victim's death, whereas the latter merely ends her conscious mental activity until she dies in a causally unrelated way. Thus, how we should conceptually distinguish the two kinds of act will depend on how we define death. Though many definitions of death (and, by extension, killing) have been proposed, all fall into one or the other of two categories. The first category, which we can call the *Existence Conception* of death, identifies a person's death with the permanent extinction of what she essentially is. On this view, the metaphysical difference between killing someone and knocking her unconscious until death is that while the former ends her existence, the latter merely causes her to have no conscious mental activity until she ceases to exist.

The second category, which we can call the *Biological Conception* of death, identifies a person's death with the end of the functioning of one or more of her basic organismic processes. This conception includes the traditional cardiopulmonary criterion and the now-popular whole brain criterion.<sup>22</sup> On the Biological Conception, the difference between killing a person and knocking her unconscious until death is that while the former causes the permanent cessation of the relevant biological processes, the latter merely causes her to have no conscious mental activity up to the time those processes cease from independent causes.

These two conceptions – Existence and Biological – cover all plausible definitions of death. Hence, they cover all plausible views about what distinguishes killing from knocking unconscious until death. Therefore, if it can be shown that on neither conception of death does the fact that one causes death constitute a wrong-making property of an act, we will have shown that the two types of act are morally equivalent.

I begin with the Existence Conception of death. On this conception, killing ends someone's existence while knocking someone unconscious until death merely causes her conscious mental life to cease up to the time that her existence ends for other reasons. Hanser seems to appeal to this understanding of killing when he writes that the reason killing a person is more wrong than knocking her unconscious is that '[s]omeone who has been ... knocked unconscious continues on in an impaired state; someone who has been killed does not continue on at all'.<sup>23</sup>

This point recalls one made by Frances Kamm. She writes:

<sup>21</sup>McMahan, *The Ethics of Killing*, pp. 241–3.

<sup>22</sup>See e.g. James Bernat, 'A Defense of the Whole-Brain Concept of Death', *The Hastings Center Report* 28 (1998), pp. 14–23.

<sup>23</sup>Hanser, 'The Wrongness of Killing', p. 399.

Suppose we put someone into a coma, knowing that he will never recover from it ... In such a case, we would not, I believe, be treated as harshly as if we had killed someone ... The fact that one person determines the nonexistence of another against his will ... is a factor in making killing wrong.<sup>24</sup>

If Kamm and Hanser are right that it is in itself morally significant whether an act determines another person's non-existence, then the Existence Conception of death would seem to imply the non-equivalence of killing and knocking unconscious until death.

Whether my action ends a person's existence – and hence whether it counts as a killing on the Existence Conception of death – depends on what the condition(s) are for a person to continue to exist. I shall argue that on none of the mainstream accounts of what it takes for a person to continue to exist is there a moral difference between killing her and knocking her unconscious until death.

Consider first a psychological continuity view, according to which a person will exist in the future if and only if there is some individual to whom she is appropriately connected via overlapping beliefs, memories, desires that lead to actions, and so on.<sup>25</sup> Because most of our psychology can exist without being consciously active, on a plausible understanding of the psychological continuity condition it will be possible to knock someone unconscious until her death without causing her to cease to exist. On this condition for continued existence, the difference between killing someone and knocking her unconscious until death is that the latter does not cause the non-existence of the psychological states that ground her persistence.

A second view is what Jeff McMahan calls the 'embodied mind account'.<sup>26</sup> According to this view, a person continues to exist if and only if there is the physical and functional (or potentially functional) continuity of enough of those parts of the brain to retain the capacity for consciousness. On the embodied mind account, then, the difference between killing someone and merely knocking her unconscious until her death is that only the killing involves the destruction of enough of the relevant parts of the brain.

I believe that neither on the psychological continuity condition, nor on the embodied mind account, is it morally relevant whether one's action causes someone to cease to exist. To see why, it will be instructive to consider two real ways a human being can be biologically alive even though she will never again be conscious.<sup>27</sup> The first is a permanent vegetative state, in which the victim has permanently lost consciousness because of irreversible damage to her cortex and limbic system. In a permanent vegetative state, the only parts of the brain that retain any functional capacity are the brainstem's vegetative centres: those responsible for heart rate, blood pressure control, temperature control, and respiration. The other way involves damage to the ascending

<sup>24</sup>Kamm, *Morality, Mortality*, p. 21.

<sup>25</sup>Sydney Shoemaker, 'Persons and their Pasts', *American Philosophical Quarterly* 7 (1970), pp. 269–85; John Perry, 'Can the Self Divide?', *Journal of Philosophy* 69 (1972), pp. 463–88; David Lewis, 'Survival and Identity', *The Identities of Persons*, ed. Amélie Rorty (Berkeley, 1976), pp. 17–40; Derek Parfit, *Reasons and Persons* (Oxford, 1984).

<sup>26</sup>McMahan, *The Ethics of Killing*. Derek Parfit defended a similar view in 'We Are Not Human Beings', *Philosophy* 87 (2012), pp. 5–28.

<sup>27</sup>For illuminating surveys of these phenomena, see Malcolm Horne, 'Are People in a Persistent Vegetative State Conscious?', *Monash Bioethics Review* 28 (2009), pp. 1–12, and McMahan, *The Ethics of Killing*, pp. 423–55.



reticular activating system, a network of cells in the brainstem that is necessary for any consciousness to occur.<sup>28</sup> A patient with localized damage to the reticular activating system retains the *capacity* for consciousness, as well as the bulk of her distinctive psychology, which is stored in or constituted by parts of the cerebrum. But because of the damage, as a matter of fact she will never again have any conscious experience.

Suppose that the conditions for an individual to continue to exist are either the continued existence of her psychology or the physical and functional continuity of those parts of the brain that have the capacity for consciousness. If it were in itself morally objectionable to terminate a person's existence, then on either of these accounts it would be objectionable to end a patient's life by terminating life support if she were in a total coma from which she would never recover due to damage to the reticular activating system, but not to remove life support from a patient in a permanent vegetative state. That is because the former ends the continuity of the patient's psychology as well as the capacity for consciousness, whereas in a permanent vegetative state they have already been destroyed. It is very hard to believe that there is this difference.

The foregoing argument shows, I believe, that on neither of these two accounts of what it takes for someone to continue to exist is it in itself morally objectionable to cause a person's non-existence. But it shows this by appealing to atypical cases in which the conscious life of the victim has already ended forever. It might be replied that terminating a person's existence *is* especially objectionable, but only so long as the conscious life of the victim is not already over. On this view, the wrongness of ending a person's existence is a 'combination effect': when ending a person's existence also ends her conscious life, the former constitutes an additional moral objection over and above the wrongness of the latter. It would not, however, be more wrongful to cause a conscious person to be in a permanent vegetative state than it would be to cause her to be permanently unconscious by destroying her reticular activating system. That is so even though only the former would involve the termination of psychological continuity as well as the destruction of those parts of the brain that have the capacity for consciousness. It is not plausible that there is a stronger moral objection to ending someone's conscious life forever based only on which part of the brain one damages. And although the psychology of a patient with permanent damage to her reticular activating system survives, practically speaking it is only 'stored'. The mere fact that a person's non-conscious psychology continues to exist could not, I believe, mitigate the wrongness of forever terminating her conscious life.

Lastly, consider a biological continuity condition of existence, according to which a person will exist in the future if and only if her organism continues to exist and function biologically.<sup>29</sup> If continued existence is conditional upon biological continuity, however, then being killed on the Existence Conception coincides with being killed on the Biological Conception. Because I will also argue that killing someone and knocking her unconscious until death are morally equivalent when killing is understood on the Biological Conception, I can argue that they are morally equivalent on the Existence Conception incidentally, since my arguments will also apply to it.

Let us turn, then, to the Biological Conception of death, according to which the difference between killing and knocking unconscious until death is that the former but not the latter involves the termination of the victim's vital biological processes. In some

<sup>28</sup>Horne, 'Persistent Vegetative State'; Adam Zeman, 'Consciousness', *Brain* 124 (2001), pp. 1263–89.

<sup>29</sup>Eric Olson, *The Human Animal: Personal Identity without Psychology*, (Oxford, 1997); Paul Snowdon, *Persons, Animals, Ourselves* (Oxford, 2014).

passages, Hanser seems to be adopting this conception of killing instead of the Existence Conception. He writes:

An agent arguably interferes with someone in a much more fundamental way by killing him than he does by knocking him unconscious ... The former causes the complete and permanent cessation of the victim's basic life-sustaining (or life-constituting) bodily operations; the latter leaves the victim alive but, for a time, unable to exercise a variety of his practical and perceptual capacities.<sup>30</sup>

Hanser here cites two distinct properties of killing someone which allegedly make it more wrong than knocking someone unconscious. One is that it causes the cessation of the victim's vital bodily operations; the other is the permanence of that cessation. These components are separable. Medical advances might one day permit the biological resuscitation of a body that has ceased to function completely.

We should ask which of the two properties (or both) is morally significant. Consider first the appeal to killing's *permanence*. Assuming that one's victim is mortal, knocking her unconscious does not change the fact that there is *a* time at which she will be permanently dead. So, if killing a person is more objectionable than knocking her unconscious because of the permanence of death, that must be because by killing someone one is the *cause* of her permanent biological death. To see whether this is a wrong-making property of killing, contrast a pair of cases in which, without any interference, the victim would have lived forever:

- (1) I cause my victim's vital biological operations to cease permanently at  $t_1$ , thereby pre-empting their permanent cessation at  $t_2$ .
- (2) I cause my victim's vital biological operations to cease for the interval  $t_1$ – $t_2$ . At  $t_2$ , they will be made to cease *permanently* by some independent process (one I fully foresee but do not cause). That other process would have occurred when it did whether or not I had interfered with my victim first.

It seems to me that what I do in (1) is no more wrong than what I do in (2), even though in only the former case do I cause the victim's permanent biological death. Causing the temporary cessation of someone's biological functions, which will foreseeably cease permanently and immediately thereafter, is no less wrong than causing their permanent cessation. In support of this claim, consider an analogous pair of cases in which the harm for which I am responsible is non-fatal:

- (3) Albert will be infected by the incurable debilitating Virus X next week if and only if I don't pre-emptively infect him with it today.
- (4) Bernard will be infected by Virus Y next week whether or not I infect him with Virus X today. Virus Y has identical symptoms to Virus X, but it will also completely inhibit Virus X if that virus is present in the host when Virus Y is contracted.

What I do in (3) is not, I believe, more seriously objectionable than what I do in (4). Some support for this claim comes from the observation that if I had to infect just one of these men, I would have no more reason to choose Bernard. That is so even

<sup>30</sup>Hanser, 'The Wrongness of Killing', p. 399.

though only in Albert's case would I be responsible for the event that is someone becoming debilitatingly ill for the rest of his life. Suppose, furthermore, that we reimagine the cases so that the viruses are beneficial, rather than harmful, to their hosts. I would have no more reason to give Albert the beneficial Virus X\* than I would to give X\* to Bernard, who will contract the beneficial Y\* next week no matter what I do. Assuming that benefits and harms are analogous in the relevant way, this supports the claim that, in instances of non-lethal harming, the wrongness of one's act is determined by the difference one makes to what happens to the victim, not by the effect on her for which one is responsible. Without a good reason not to, it seems to me that we should generalize this result to lethal harms. The mere fact that killing someone causes the permanent cessation of her vital biological operations is not, I conclude, a factor in making it wrong.

Instead of appealing to killing's permanence, one might try to show that killing is especially wrong by appealing to its *irreversibility*. It might be said that depriving someone of conscious life by killing her is more objectionable than doing so by knocking her unconscious because the effect of the former on the victim cannot be undone. We can understand this claim to be about metaphysical, nomological, or practical irreversibility. Consider first the proposal that killing is especially wrongful because it is either metaphysically or nomologically irreversible. Suppose I wilfully terminate all of a person's mental processes along with her biological life. Her condition is not metaphysically or nomologically irreversible, for I know that a drug that would revitalize her has been made in the past. However, I also know that all samples of the drug have been destroyed and the knowledge of how to manufacture it has been lost forever. It is both metaphysically and nomologically possible to make the drug again, so my victim's condition is reversible in those senses; I just know that it will not be reversed. I believe that it would be no less wrong to end my victim's life in this case than it would be had the drug never existed and indeed never could exist.

Consider next the claim that killing is especially wrong because death is irreversible in a *practical* sense; killing a person, we might say, is objectionable in part because it makes it so that there is no feasible way to bring that person back to life. Against this claim, suppose the aforementioned drug exists, but the last remaining sample is owned by an independently acting agent who I know will never use it to restore the life of my would-be victim. Assume that I know, too, that there is nothing I can do to change this fact. When I end my victim's life in this case my effect is reversible in the practical sense because there is a feasible way to bring her back to life. Yet it seems to me that it would be no less wrongful to end her life in that case than it would be if the drug had never existed. Because this argument appeals to what I and another agent do, it may seem misleading to focus on the wrongness of my action alone. But the mere fact that there is someone *else* who could revive my victim does not mitigate the wrongness of my own action, given that I know, when I act, that this other person will never in fact revive her and cannot be made to do so. Ending her life would be no less wrongful if the reason the drug could not be used was some impassable natural obstacle instead of an agential one.

I have argued that neither death's permanence nor its irreversibility is a factor in making killing a person more wrong than knocking her unconscious until death. Let us turn, then, to the more fundamental wrong-making feature of the Biological Conception: simply that it brings about the cessation of a person's life-sustaining (or life-constituting) bodily operations. I believe that the fact that one act causes the cessation of a person's bodily operations does not, on its own, make its performance more

objectionable than another. Suppose that someone in a temporary and medically non-serious coma will soon be connected to a life support system. While she is connected to the machine, if the parts of her brain that support her mental life are destroyed, her bodily operations will continue to function: her heart will pump blood, her lungs will inhale and exhale, and so forth. If those parts of her brain are destroyed while she is not connected to the machine, the injury will cause her biological death. If bringing about the cessation of a person's vital biological operations were in itself wrongful, then it would be less objectionable to obliterate this person's brain while she was on the life support system than to do so before she was connected to it, for in the former case her biological operations would continue to function even after the totality of her mental life was gone. But it is very difficult to believe that these two acts differ in their degrees of wrongness.

Earlier in this section, I postponed arguing that it is not in itself wrongful to terminate a person's existence if personal persistence is a matter of biological continuity. I take the foregoing arguments to have established that conclusion. Hence, on none of the mainstream accounts of the conditions for our continuity is terminating someone's existence a factor in making killing wrong.

To summarize, on neither the Existence Conception nor the Biological Conception of killing is there is an intrinsic moral difference between killing a person and knocking her unconscious until her death. Because these two understandings exhaust the plausible conceptions of killing, and because one type of act is more wrong than another only if there is some morally relevant essential property of the former that is not shared by the latter, I conclude that it is just as wrong to knock someone unconscious until her death as it is to deprive her of the same amount of conscious life by killing her.

#### V. The temporal irrelevance of knocking unconscious

I turn now to the second premise in my argument against the Equal Wrongness Thesis:

- (B) It is as wrong to deprive someone of a certain amount of conscious life by knocking her temporarily unconscious in the middle of her life as it is to deprive someone of the same amount of conscious life by knocking her unconscious until her death, other things being the same.

Before proceeding, it is worth clarifying two things about (B)'s *ceteris paribus* clause. First, it should be understood to exclude any differences in the prudential value for the victim of the conscious period of which she is deprived. It may be said that the last month of a person's life would have been especially important because it would have allowed her to say her goodbyes, or to realize some end towards which she had been working. But these are contingent facts about a life, and anyway I suspect that for many people who will live until old age a month in the middle of their life is worth at least as much to them as one at its end. Second, (B) should be understood to exclude differences in the length of the victims' conscious lifespans. It is plausible that it is more wrong to deprive someone of a given amount of consciousness the shorter her total conscious life would have been, but the premise applies only when that difference is factored out.

To get a better sense of (B), consider the following illustration. Calvin and Dylan will both live for exactly 1,000 months (about 83 years). Calvin has lived for 500 months so far while Dylan has lived for 999 months. Apart from their ages, the two men differ in

no morally relevant respect. Suppose that the prudential value of what would be Dylan's 1,000th conscious month is the same as that of what would be Calvin's 501st. (B) says that the strength of the moral objection to knocking out either man for one month is the same.

There are several ways to object to this claim. One is to argue that the deprivation of a month of conscious life is *worse* for Dylan than it is for Calvin. Because we are assuming that the prudential value of the month each man loses is the same, this cannot be a claim about lifetime well-being. Although it is true that Dylan has less remaining life, he has to that extent lived more life already.

It might be claimed instead that the deprivation of his last conscious month constitutes a greater loss for Dylan in the *non-comparative* sense. Frances Kamm has argued that death is a non-comparatively bad event because it involves the loss of one's last period of conscious experience and therefore means that 'everything is all over for one'.<sup>31</sup> If Kamm is right, then even if a conscious period at the end of a person's life would have contained no more value for him than one in the middle of his life, depriving him of the final period would be more harmful because it involves an additional, non-comparative harm. That might make knocking someone unconscious until his death more wrongful than knocking him temporarily unconscious.

If the permanent end of consciousness really is a non-comparatively bad event, we should in principle be able to ask how bad it is. In the first place, however, it is unclear how to go about assigning it a non-arbitrary value such that it is equivalent to some amount of pain, for example. But let us grant that we can. It is still implausible that a person's death, or the permanent end of her consciousness generally, constitutes a non-comparative bad for her. Suppose, just for the sake of exposition, that that non-comparative badness is such that to offset its presence in a person's life she would need to live for at least four happy years. Then a life of three happy years would not be worth living, a life of two worse still, and so on. But I cannot believe that the life of a child who dies after two happy years is overall bad. Most people would be glad for the child's own sake that it got to live at all. We might judge its death a tragedy because human beings typically live for much longer. But the lifespan of typical members of one's species does not seem normatively important.

Of course, perhaps the permanent end of consciousness is just not as bad as all that. Let us suppose it is only bad enough to outweigh three happy months. Would anyone really conclude, upon the discovery of an extra-terrestrial species whose members die painlessly after only one happy summer, that the total well-being in the universe is lower than previously thought? Defenders of the position that the final end of consciousness is non-comparatively bad for its subject might reply that the position applies to persons only. But things seem no different if we clarify that these extra-terrestrials spring into being as fully formed persons. For my part, it is hard to believe that the permanent end of consciousness could be non-comparatively bad enough to outweigh even one happy day.

Instead of appealing to the non-comparative badness of being deprived of one's last remaining month of consciousness, we might say that depriving someone of that month is especially wrongful because it is worse for him in a *time-relative* sense. Other things being equal, we might think it is time-relatively worse to lose a given amount of future good, the less good he has to look forward to. Even if time-relative interests should be given special weight in one's decisions concerning one's own life, it is not, however,

<sup>31</sup>Kamm, *Mortality, Mortality*; Kamm, 'The Purpose of My Death', pp. 734–5.

clear that they should guide one's decisions concerning the treatment of others. Suppose I will not hear from a friend for several months.<sup>32</sup> I know the following: at some point during that time she will either undergo an excruciatingly painful operation, or else a lesser-but-still-seriously painful operation. Suppose I learn today that either she has already undergone the more painful operation or else she will undergo the less painful one next week. Relative to today, it is (we are assuming) in her interest that she has already suffered the greater pain. However, it does not seem that *I* should take this position. I should hope, for my friend's sake, that she will undergo the less painful but future episode, because then her life will contain less pain overall. Indeed, it seems to me that I should view the situation no differently from the way I would in a case in which both potential pains are in the future. This seems to indicate that in so far as the moral salience of a decision derives from its potential to affect the well-being of others, it is the lifetime perspective that matters. Although there may be a sense in which an unconscious episode at the end of one's life is a greater loss, it is not the sense that should lead us to doubt (B).

We cannot, I conclude, appeal to the special badness of losing one's last period of consciousness in order to refute (B). The other way to challenge the premise is to skirt questions of badness and proceed directly to the wrongness of the acts. In a response to Lippert-Rasmussen's article, Daniel Cohen and Morgan Luck have argued that the wrongness of depriving a victim of a given amount of consciousness is greater the nearer the victim is to the end of her conscious life.<sup>33</sup> More precisely, the authors defend the following principle:

The wrongness of rendering someone unconscious corresponds with the proportion of the victim's *remaining* conscious life that is thereby removed, other things being the same.<sup>34</sup>

Their argument for that principle involves imagining two alien species. Members of the first species (call them the 'short-livers') always live for just two years, while members of the second species (call them the 'long-livers') live for 1,000 years. The authors imagine coming across a newborn member of each species.<sup>35</sup> It seems more seriously wrongful to render the newborn short-liver unconscious for one year – half its remaining life – than to knock the newborn long-liver unconscious for a year, leaving it to experience another 999 years when it awakens.

Cohen and Luck's principle gets the right answer in this case, since the short-liver is deprived of a greater proportion of its remaining conscious life (one year of two) than the long-liver is of his (one year of a thousand). But the implications of their principle are at least as counterintuitive. Suppose, instead, that we must knock unconscious for a year either one of their short-livers who has been alive for one year, or else a long-liver who has already lived for 999 years. The principle Cohen and Luck propose implies that these acts are equally wrongful, for in either case one removes the entirety of the victim's remaining conscious life. When we knock the short-liver unconscious we deprive

<sup>32</sup>The argument that follows draws on one made by Parfit (*Reasons and Persons*, ch. 8) and an example by Caspar Hare ('A Puzzle about Other-Directed Time-Bias', *Australasian Journal of Philosophy* 86 (2008), pp. 269–77).

<sup>33</sup>Cohen and Luck, 'Why a Victim's Age Is Irrelevant', pp. 396–401.

<sup>34</sup>Cohen and Luck, 'Why a Victim's Age Is Irrelevant', p. 399 (emphasis mine).

<sup>35</sup>Assume that members of both species are born full persons.

it of half its total conscious life, whereas when we do so to a long-liver we reduce its conscious life by just one-thousandth.

It might be thought that the *ceteris paribus* clause of Cohen and Luck's principle is meant to factor out differences in how much conscious life the victim has so far enjoyed. But their conclusion is that a person's age has no bearing on the wrongness of killing her, so their argument cannot go through if cases in which the victims differ in the amount of conscious life they have lived are not covered by their principle.

Cohen and Luck's newborn aliens case relies on the intuition that it matters, when knocking a person unconscious, how much of her future experience she loses thereby. But it also matters how much conscious experience a person has so far enjoyed. To account for both factors requires a middle ground. One plausible candidate is that the wrongness of knocking a person unconscious for some period corresponds not with that period's proportion to the victim's remaining conscious life, but rather with its proportion to the total lifetime conscious experience she would otherwise have had.

That new principle would be:

The wrongness of knocking someone unconscious corresponds with the proportion she loses of the *total* amount of conscious life she would otherwise have enjoyed, other things being the same.

This principle gets the right answer in both of the two short- and long-lived alien cases. It also entails (B), which compares equally long deprivations of consciousness at different times within equally long conscious lives.

I now turn to a positive argument for (B). In it, Calvin and Dylan again make an appearance. Suppose that I possess two drugs: a *fast-acting drug* that immediately knocks its taker unconscious for one month and a *delayed-acting drug* that has no effect until one month before its taker's death, when it will knock him unconscious for that final month. Consider, first, the following case:

*Case One:* I secretly slip the fast-acting drug into the coffee of the 999-month-old Dylan and the delayed-acting drug into the coffee of the 500-month-old Calvin.

It is not plausible that what I do to either man in Case One is more wrong than what I do to the other. I deprive each man of the very same amount of experience, at the very same location in his life, and in the very same manner. It is true that the effect of my action takes place immediately for Dylan but is delayed by nearly forty years for Calvin. But it seems to me that mere temporal distance between an act and its effect could not be morally significant in itself.

Next, consider:

*Case Two:* Calvin has an identical twin, Caleb. I slip the delayed-acting drug into Calvin's coffee, and the fast-acting drug into Caleb's.

I do not think it is plausible that the acts in Case Two are wrong to different degrees. To dispute that claim, one cannot appeal to the wrongness of making it the case that 'everything is all over' for him, since both victims have conscious life in their future. Moreover, both men are deprived of the same proportion of their *remaining* conscious lives. The salient difference seems to be that only Calvin loses the final month of his

conscious life. But since it is not the last month he has left, and since there is not a morally relevant sense in which being deprived of that month of consciousness constitutes a greater loss to him, it is hard to see why the mere fact that the month is at the end of his life should make its deprivation more wrongful.

Taken together, these two equivalences – that in Case One and that in Case Two – show that it is morally irrelevant whether, when I deprive someone of a month of conscious experience, I do so at the end of his life or its middle. If giving the fast-acting drug to the 999-month-old Dylan is morally equivalent to giving the delayed-acting drug to the 500-month-old Calvin (Case One), and the latter act is morally equivalent to giving the fast-acting drug to the 500-month-old Caleb (Case Two), then giving the fast-acting drug to the 999-month-old Dylan is morally equivalent to giving the 500-month-old Caleb the fast-acting drug. That result, of course, is the same as the case with which we began: knocking out Calvin now, or knocking out Dylan now.<sup>36</sup>

## VI. Conclusion

It is now possible to bring together the pieces of my argument. The first two premises were:

- (A) It is as wrong to deprive someone of a certain amount of conscious life by knocking her unconscious until her death as it is to deprive someone of the same amount of conscious life by killing her, other things being the same.

And

- (B) It is as wrong to deprive someone of a certain amount of conscious life by knocking her temporarily unconscious in the middle of her life as it is to deprive someone of the same amount of conscious life by knocking her unconscious until her death, other things being the same.

These two premises entail:

- (C) It is as wrong to deprive someone of a certain amount of conscious life by knocking her temporarily unconscious in the middle of her life as it is to deprive her of the same amount of conscious life by killing her, other things being the same.

(C) is sufficient to defeat the Equal Wrongness Thesis. Suppose we take the period of unconsciousness to be one month. We would then hold that knocking someone unconscious for one month in the middle of her life is as wrong as killing someone who has one month left to live. The Equal Wrongness Thesis tells us that killing a person who has one month left to live is as wrong as killing a person with fifty good years to live. Hence, the conjunction of the Equal Wrongness Thesis and (C) entails that, other things being equal, knocking a person unconscious for a month in the middle of her life is just as wrong as killing a person with fifty good years ahead of her. This result is absurd. Whatever considerations might justify knocking someone out for a month in the middle of her life, those considerations cannot be as strong as the ones required

<sup>36</sup>I believe that my arguments generalize to longer deprivations of consciousness than one month. As we will see presently, however, even that specific equivalence is sufficient to undermine the Equal Wrongness Thesis.



to justify killing a person with fifty good years left to live. On pain of absurdity, therefore, we should reject the Equal Wrongness Thesis.<sup>37</sup>

Author ORCID:  Todd Karhu, [0000-0003-1726-7257](https://orcid.org/0000-0003-1726-7257)

---

<sup>37</sup>For valuable feedback on this article I am grateful to Susanne Burri, Goreti Faria, Tomi Francis, David Kinney, Kasper Lippert-Rasmussen, Christopher Marshall, Michal Masny, Jeff McMahan, Max Muir, Michael Otsuka, Bryan Roberts, Tom Rowe, Bastian Steuwer, Alex Voorhoeve, two anonymous *Utilitas* referees, and audiences at the London School of Economics and the University of Reading.

**Cite this article:** Karhu T (2019). Not All Killings Are Equally Wrong. *Utilitas* 31, 378–394. <https://doi.org/10.1017/S0953820818000328>