

Article

# The Meat Eater Problem

Michael Plant

Wellbeing Research Centre, University of Oxford and the Happier Lives Institute;  
michael.plant@hmc.ox.ac.uk

Submitted: 18 February 2021, accepted: 9 September 2022, published: 31 October 2022

**Abstract:** Here are two commonly held moral views. First, we must save strangers' lives, at least if we can do so easily: you would be required to rescue a child drowning in a pond even if it will ruin your expensive suit. Second, it is wrong to eat meat because of the suffering caused to animals in factory farms. Many accept both simultaneously—Peter Singer is the pre-eminent example. I point out that these two beliefs are in a sharp and seemingly unrecognised tension and may even be incompatible. It seems universally accepted that doing or allowing a harm is permissible—and may even be required—when it is the lesser evil. I argue that, if meat eating is wrong on animal suffering grounds then, once we consider how much suffering might occur, it starts to seem plausible that saving strangers would be the greater evil than not rescuing them and is, therefore, not required after all. Given the uncertainties and subjective assessments here, reasonable people could substantially disagree. The surprising result is that a moral principle widely considered to be obviously true—we must rescue others—is not, on further reflection, obviously true and would be defensibly rejected by some. Some potential implications are discussed.

**Keywords:** applied ethics; easy rescue; Shallow Pond; animals; Peter Singer

**How to cite:** Plant, M. The Meat Eater Problem. *Journal of Controversial Ideas* 2022, 2(2), 2; doi:[10.35995/jci02020002](https://doi.org/10.35995/jci02020002).

©2022 Copyright by the author. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license.



## 1. Introduction

It is widely believed that we, as members of the public, have a Duty of Easy Rescue to one another.

*Duty of Easy Rescue:* We are required to save lives in *rescue cases*, one-off instances where we can physically save a stranger at trivial cost to ourselves.<sup>1</sup>

---

<sup>1</sup> The singularity and physicality of the specification are deliberate, as the intuition is strongest that we must save the person in this case, and the point is to test the strongest version. Intuitions are weaker if there

To illustrate this, consider the following, familiar case from Singer<sup>2</sup>:

*Shallow Pond*: You are walking past a shallow pond and see a drowning child. You can easily rescue the child, but doing so will ruin the expensive new suit you are wearing.

Intuitively, we are required to save the child. This is because, as Peter Singer explains: “[it] will mean getting my clothes muddy, but this is insignificant, while the death of the child would presumably be a very bad thing.”<sup>3</sup>

It is also widely believed that it is wrong to be a *meat eater*, someone who regularly consumes animal products produced from factory farms.<sup>4</sup> This is on the grounds that this consumption requires creating animals who, due to the conditions in factory farms, live overall bad lives.<sup>5</sup>

While there are other reasons one might believe that meat eating is wrong—the environmental impact, the violation of animals’ rights, etc.—here I only focus on those arising from animal suffering caused by factory farming as it is the easiest way to generate the problem I shortly raise.

Many people hold both beliefs and, what’s more, do not consider this conjunction problematic. The most obvious example is Peter Singer: he is the originator of the Shallow Pond case and encourages people to give to life-saving charities;<sup>6</sup> he has publicly campaigned against factory farming on animal suffering grounds and advocated veganism for many years.<sup>7</sup> The simultaneous acceptance of these two beliefs is common among moral philosophers and, increasingly, in society at large.

This paper argues these two beliefs are, in reality, in substantial tension and may well be incompatible, once additional plausible empirical and normative considerations are accounted for. Further, if they are incompatible, we must abandon the notion that there is a Duty of Easy Rescue.

Here, in brief, is the argument. We accept that doing (or allowing) harm is permissible—and may be required—when it is the lesser evil; to put the same thing differently, we are not required to do (or allow) the greater evil. Therefore, we wouldn’t be required to save lives, even if we could do so easily, if that would be the greater evil. I argue that, if we accept that meat eating is wrong (on animal suffering grounds) then, once we look at the details, it starts to look plausible that meat eating causes so much suffering that saving the lives of strangers would be the greater evil (compared to not

---

are hundreds of people we could save, each time at a cost to our self, as that may be too demanding: Timmerman, T. (2015). “Sometimes there is nothing wrong with letting a child drown,” *Analysis* 75(2): 204–12.

<sup>2</sup> Singer, P. (1972). “Famine, affluence, and morality,” *Philosophy & Public Affairs* 1(3): 229–43, p. 231.

<sup>3</sup> Ibid

<sup>4</sup> I will use the term ‘meat eaters’ as a shorthand for ‘factory-farm-produced-animal-product-consumers’. This is admittedly imprecise—one could be the former without being the latter, and vice versa—but unproblematic in practice, given the overlap between these categories, as least in high income countries: the vast majority of people eat meat and the vast majority of meat is produced in factory farms. See footnote 45 for more.

<sup>5</sup> Note that, on this view, consumption is only problematic because and to the extent that it creates demand for further production of overall unhappy lives.

<sup>6</sup> Singer (1972); Singer, P. (2015). *The Most Good You Can Do: How Effective Altruism Is Changing Ideas about Living Ethically*. Text Publishing.

<sup>7</sup> Singer, P. (1975). *Animal Liberation, Animal Liberation*. London: Jonathan Cape.

saving them) and would, therefore, not be required. Simply, accounting for the existing concern for animal welfare reduces, and may remove, the obligation to rescue others.

The most controversial part of the argument is, of course, that meat eating causes so much suffering that saving strangers is the lesser evil. I expect readers will consider that preposterous.

Let's consider each of the four steps of the argument in slightly greater detail now. We will then spend the bulk of the essay exploring the controversial one.

The first move is to note that there is widespread (perhaps universal) agreement that there can be *lesser evil* justifications for doing or allowing harm. For consequentialists, we are required to choose the lesser evil, that is, the better outcome. For non-consequentialists, it is at least permissible to do/(allow) some harm to prevent a greater harm.<sup>8</sup> The classic case illustrating this is the Trolley Problem:<sup>9</sup>

*Trolley Problem:* A trolley is rolling down a track and will kill five people. You can pull a lever so that the trolley switches onto another track and kills one person instead.

The common intuition is that we are at least permitted to pull the lever, even though that will kill an innocent person, because this is the lesser evil.

Many non-consequentialists believe there is an important distinction between doing and allowing harm: there is a normative badness associated with doing harm, which means that you are only permitted to do harm to prevent another harm if the harm you prevent is *substantially* greater.<sup>10</sup> Thus, it might be permissible to kill one to save five, but not to kill one to save two. In this essay, the doing-allowing distinction is unimportant as we will only be concerned with a case of allowing harm to prevent further harm. A further issue for non-consequentialists, and one I will not discuss, is whether one is sometimes required to bring about the lesser evil—in this case, to pull the lever.<sup>11</sup> It is worth emphasizing that both consequentialists and non-consequentialists agree that, at the least, we are *not required* to bring about the greater evil.<sup>12</sup>

The second step is observing that we would not be required to save someone's life if doing so were the greater evil. Here is a motivating case:

*Drowning Dictator:* You live in a country ruled by a dictator. One day, when walking past a pond you see him drowning—he has a distinctive appearance, and you recognise him from his ubiquitous propaganda posters. You can easily rescue him. However, you realise that saving him will not only ruin your new expensive shoes, but he will foreseeably go on to torture and terrorise thousands of people in the future.

---

<sup>8</sup> When I say 'at least permissible', I mean permissible or required, that is, not impermissible. Feinberg, J. (1978). "Voluntary euthanasia and the inalienable right to life," *Philosophy & Public Affairs* 7(2): 93–123; Thomson, J. J. (1986). "Some ruminations on rights," in *Rights, Restitution, and Risk: Essays, in Moral Theory*. Cambridge, MA: Harvard University Press, pp. 49–65.

<sup>9</sup> Thomson, J. J. (1976). "Killing, letting die, and the trolley problem," *The Monist* 59(2): 204–17.

<sup>10</sup> McMahan, J. (2009). *Killing in War*. Oxford University Press, p. 29.

<sup>11</sup> Walen, A. and Wasserman, D. (2012). "Agents, impartiality, and the priority of claims over duties: Diagnosing why Thomson still gets the Trolley Problem wrong by appeal to the 'mechanics of claims'," *Journal of Moral Philosophy* 9(4): 545–71; Frowe, H. (2018). "Lesser-evil justifications for harming: Why we're required to turn the trolley," *The Philosophical Quarterly* 68(272): 460–80.

<sup>12</sup> Note that lesser evil justifications can be insensitive to concerns about moral responsibility. Hence—and while I cannot think of a plausible reason for this—if not saving a vegan did happen to have a big animal-suffering-reducing effect, then that would be the lesser evil and so not required either; I am grateful to Jeff McMahan for pointing this out.

Clearly, letting the dictator drown is the lesser evil and would therefore be at least permissible. Note that Drowning Dictator is not a case of *causing* a harm to prevent a greater evil; rather, it seems akin to a choice between allowing a lesser evil—the dictator’s death—and allowing another, greater evil—the suffering that results if the dictator lives. As noted, there seems to be consensus that we are not *required* to allow the greater evil.

Indeed, we might think that we would not be required to save the dictator, even if the harm of him of dying was somewhat larger than the harm he would cause to others: if we save him, that might make us complicit, in some morally problematically sense, in his crimes. While I mention this, I won’t develop this point further.

Besides the lesser evil justification, there is also, in parallel, a *liability* justification for not saving the dictator, much as we think an assailant attacking an innocent violent is liable to be harmed to protect the victim. There is even a strong case that the dictator would be liable to be killed, rather than merely not rescued. Similarly, we might wonder if meat eaters are liable to be harmed to protect the animals.

In this essay, I will only explore the lesser evil justification.<sup>13</sup> There are several reasons for this. First, it’s not at all clear that there would be an equivalent liability justification for meat eating. One might argue someone can only be liable to be harmed if they are *culpable*<sup>14</sup> and that meat eaters are not culpable for their actions—they are simply doing what everyone else does.<sup>15</sup> Further, it’s unclear if we can violate an entity’s rights by creating it when it could not otherwise have existed. Second, liability to harm may, in any case, depend on whether that harm is *proportional*,<sup>16</sup> in which case we need to examine how much harm meat eating causes anyway. Third, both consequentialists and non-consequentialists accept the lesser evil justification, whereas the former do not accept liability-based ones; hence the lesser evil argument has general application.

To be clear, the Drowning Dictator case poses no threat to the Duty of Easy Rescue, at least as I have specified it above. This is because the Duty refers to saving *strangers*—people whose identity we do not know. Framing the Duty of Easy Rescue this way seems appropriate: many think we should save the lives of others we don’t know, at least if this is easy.<sup>17</sup> The Dictator case is a rare exception to the rule. We both recognise the person and reasonably expect their continued existence will do much more harm than

---

<sup>13</sup> If we wanted a similar case where there was only a lesser evil justification, perhaps we could imagine a fanciful scenario where the dictator was doing bad things only because he was intoxicated by some powerful drug, or under a spell. In this case, we might say he is causally, but not morally, responsible and thus has not forfeited his right not to be harmed. We might go on to say it would be permissible to let this ‘drugged dictator’ drown if this led to a lower total of harms whilst simultaneously denying it would be permissible to kill him. Why? Because of the standard doing–allowing concerns: he hasn’t forfeited his right not to be harmed, so to violate that right we would need to be preventing substantially larger harms compared to a case where we are allowing him to die.

<sup>14</sup> Ferzan, K. K. (2005). “Justifying self-defense,” *Law and Philosophy* 24: 711–49; Quong, J. (2012). “Liability to Defensive Harm,” *Philosophy & Public Affairs* 40(1): 45–77.

<sup>15</sup> Norcross, A. (2004). “Puppies, pigs, and people: Eating meat and marginal cases,” *Philosophical Perspectives* 18(1): 229–45; Abbate, C. E. (2020). “Meat eating and moral responsibility: Exploring the moral distinctions between meat eaters and puppy torturers,” *Utilitas* 32(4): 398–415.

<sup>16</sup> McMahan, J. (2017). “Liability, proportionality, and the number of aggressors,” in Saba Bazargan and Samuel C. Rickless (eds.), *The Ethics of War*. Oxford University Press, pp. 3–27.

<sup>17</sup> Eagle-eyed readers might object that the Duty of Easy Rescue is silent on saving people you know, such as your mother or your friends. I do not want to discuss the morality of saving people to whom you have a personal connection but who you think will cause substantial harm—are you required to save your mother from drowning if she is a serial killer? Very plausibly different intuitions and principles, such as partiality, apply here. I focus here on the more general case of saving an unidentified individual to whom you have no personal tie.

good. The Duty of Easy Rescue would, however, need to be abandoned if saving the average life turned out to do more harm than good.

The third move is to argue that, if meat eating is wrong on animal suffering grounds, then what I call the *Strong Anti-Carnism Thesis* is, in fact, surprisingly plausible:

*Strong Anti-Carnism Thesis:* Meat eaters cause a sufficient amount of animal suffering via their diets that not saving strangers' lives in rescue cases is (in expectation) the lesser evil.

This thesis combines moral and empirical premises I unpack later. Presumably, most readers will consider the Strong Anti-Carnism Thesis—I often shorten this to the 'Strong Thesis'—is wildly implausible: there is simply no way eating meat causes *that* much suffering.

The goal of the essay is to show that, when one looks at the details, the Strong Thesis is surprisingly plausible, although not certain, despite the fact that this conclusion is deeply disturbing. The basic argument is this: for each year of meat eating by a human, that creates about five years of chicken life. So, if we consider just that, and think that those animals have lives which are nearly as bad as human lives are good, then the Strong Thesis is no longer inconceivable. I consider many objections later, but none of them seem to offer a decisive reason to reject the Strong Thesis for someone who thinks meat eating is wrong and animals in factory farms suffer intensely.

Two important clarifications are in order here. First, I am not interested in discussing here whether, for whatever reason, eating meat is wrong, a subject on which there is now an enormous philosophical literature.<sup>18</sup> I do not argue that meat eating is (on suffering grounds) wrong, although I address one objection to it later. Rather the purpose of the essay is to explore what might follow *if it is the case* that meat eating is wrong and raise a challenge for those, like Singer, who think that it is.

Second, I should emphasise that I do not and cannot show the Strong Thesis is true: assessing it necessarily requires making subjective assessments and reasonable people will disagree about these. My objectives are (1) to show (at least for those who accept that meat eating is wrong) that the Strong Thesis should not be rejected out of hand, (2) it plausible enough that we can expect some to accept it, and (3) if it is true, it has at least one highly revisionary implication.

The fourth move is recognizing that it is inconsistent to accept (A) we have a Duty of Easy Rescue, (B) the Strong Thesis is true, and (C) that there are lesser evil justifications for doing or allowing harm. Let's spell out why. The Duty of Easy Rescue requires us to save strangers' lives. If the Strong Thesis is true, failing to save such individuals would be the lesser evil. If we are permitted (or required) to allow the lesser evil, then we are *not* required to cause the greater evil. Saving strangers' lives would be the greater evil and so not required. These three positions are jointly inconsistent: they entail that rescuing strangers both is, and is not, morally required. To avoid inconsistency, something must go. If the Strong Thesis were true, the only real option is to abandon the Duty of Easy Rescue. I'll say more later about why this is less counter-intuitive than it initially seems.

We are left with a surprising and unsettling conclusion. The intuition that, if someone were drowning in a shallow pond in front of you, you would be morally required to save them, is viscerally felt and seemingly beyond doubt. The puzzling implication of this

---

<sup>18</sup> See, for example, Fischer, B. and Bramble, B. (2015). *The Moral Complexities of Eating Meat*. Oxford University Press. Available at: <https://philpapers.org/rec/BRATMC-9> (accessed: 7 January 2021).

argument is that, when we account for something as seemingly unrelated as what people have for lunch, this requirement is no longer obvious and may not hold. We can expect that reasonable people would, on reflection, disagree quite substantially about how much suffering factory farming causes, with some concluding that saving lives is the greater evil. This is an unusual case where the statement ‘reasonable people will disagree’ is reasonably interesting: whether saving lives is a good thing, and we’re required to do so, are not places we would have expected disagreement.

The possibility that the harms of meat eating might make saving lives—whether in general or in rescue cases—wrong (or at, least, not required) seems to have been discussed informally in both the effective altruism and animal advocacy communities. However, I am not aware of any previous attempts—by Peter Singer or anyone else—to rigorously examine this.<sup>19</sup> The main contribution of this essay is to provide this analysis. I am not sure what, if anything, to infer from the fact the tension I consider here has not been investigated before, despite the fact Singer’s Shallow Pond case and his arguments against factory farming are almost half a century year old. I do not know if this says something general about the challenges of living and thinking ethically, or merely that this specific issue seemed too weird and too implausible to have merited attention before.

The conclusion that humans might have substantial effects through their diets raises further questions: If rescuing lives is bad, should we now conclude, further, that we are required not to rescue others? Are we now permitted to kill others? How might accounting for these effects change our priorities insofar as we strive to be effective altruists and do the most good? These issues will be briefly discussed at the end of the essay, but I do not aim to resolve them firmly there.

The remainder of this essay is structured as follows. Section 2, the longest section, sets out why, if meat eating is wrong, the Strong Anti-Carnism Thesis is surprisingly plausible. Section 3 explains why if there is inconsistency, it seems the Duty of Easy Rescue must go. Section 4 discusses some further questions the essay raises. Section 5 concludes.

## **2. If Meat Eating Is Wrong, How Plausible Is the Strong Anti-Carnism Thesis?**

Before we go any further, let’s quickly state the argument that meat eating is wrong on animal suffering grounds. I take this to be an uncontroversial explication of the view. First, animals in factory farms lead overall bad lives, that is, lives with net negative well-being. Second, the individual purchasing decisions of meat eaters do, in expectation, impact the number of animals raised and killed. Third, creating unhappy lives is bad, that is, it makes the outcome worse. Fourth, the total negative well-being caused by creating unhappy animals is sufficiently large, relative to the benefit a meat eater gets from eating meat (compared to the scenario in which they do not) to make being a meat eater wrong.<sup>20</sup>

---

<sup>19</sup> One written treatment I have found is a blog post by S. Weathers, who considers whether the idea of eating meat should cause some reduction in the value of saving lives; Weathers does not attempt to lay out the moving parts necessary to carefully make this evaluation; further, he supposes meat eating may not be bad (the created animals might be happy) and, even if it were bad, saving lives would still be good; Weathers, S. (2016). “The meat eater problem: Developing an EA response,” Effective Altruism Forum. Available at: <https://ea.greaterwrong.com/posts/gA57ThbaS3znib242/the-meat-eater-problem-developing-an-ea-response> (accessed: 5 July 2019).

<sup>20</sup> For the purposes of the argument, it’s unimportant what would count as ‘sufficiently large’.



As noted already, the main purpose of this paper is *not* to argue for this, but to highlight a further problem that arises for those who *already* accept it. I will, however, briefly discuss the second premise as it is the most controversial and readers unpersuaded by it might not engage with the rest of the paper.<sup>21</sup>

One might think that an individual's purchases make no difference to the total number of animals that get produced—the 'causal inefficacy' objection: in a complicated supply chain with farmers, wholesalers, grocers, and so on, a single purchasing decision at one end of the chain may not change the purchasing orders at all earlier steps, meaning the same number of animals live whatever each individual consumer does.

An issue with supposing that no single individual has *any* effect is that it implies, implausibly, that, if *everyone* stopped eating meat, the same number of animals would continue to be reared by farmers despite the fact no one was buying them.

We might suppose, more realistically, that there are some causal thresholds: perhaps chickens are ordered in batches of 50, so each 50th purchase results in another batch of 50. The state of the supply chain is essentially unknowable to the consumer. Therefore, each one chicken you buy has a 1/50th chance of creating 50 chickens, with the result that the expected impact of buying one chicken is that one more chicken is created and killed.<sup>22</sup> Suppliers have strong financial incentives to produce the number of animals that maximises profits and can use modern technological to accurately track what gets bought. Hence, in the ordinary case, we should consider ourselves as having a causal impact.<sup>23</sup>

How might we get from the claim that meat eating is wrong to the Strong Thesis—the claim that meat eaters cause a sufficient amount of animal suffering via their diets that not saving strangers' lives in rescue cases is the lesser evil?<sup>24</sup> There are lots of things to consider and I work through these systematically.

I start by considering just the consequences of rescuing a randomly encountered stranger compared with not saving them. This has three parts, each of which is discussed in a subsection. Section 2.1 compares the benefit to someone of having their life saved against the costs a meat-eating diet causes to animals. Section 2.2 considers the non-dietary effects the death of a meat eater would have, such as the grief caused to their friends and family. Section 2.3 accounts for the fact that not all humans are meat eaters.

For non-consequentialists, what makes something the lesser(/greater) evil does not have to depend solely on the consequences. Section 2.4 raises two reasons one might give for claiming that rescuing the stranger would be the lesser evil, even though this has worse consequences, and therefore the Duty of Easy Rescue still applies.

---

<sup>21</sup> I am grateful to Jeff McMahan for suggesting I discuss this objection.

<sup>22</sup> Singer, P. (2009). *Animal Liberation*. HarperCollins; Kagan, S. (2011). 'Do I make a difference?', *Philosophy and Public Affairs* 39(2): 105–41.

<sup>23</sup> I note some, object to this expected value argument, primarily by (re)asserting the claim there is 'slack' in the system which leads to causal inefficacy. See, for example, Budolfson, M. B. (2018). "The inefficacy objection to consequentialism and the problem with the expected consequences response," *Philosophical Studies* 176: 1711–24; and Nefsky, J. (2018). "Consumer choice and collective impact," in M. Budolfson, T. Dogget, and A. Barnhill, A. (eds), *The Oxford Handbook of Food Ethics*. New York: Oxford University Press. I find such argument unconvincing: they rely on a misunderstanding of the underlying economic realities, but to explain this here would take up too much space. For such an explanation, see McMullen, S. and Halteman, M. C. (2018). "Against Inefficacy objections: The real economic impact of individual consumer choices on animal agriculture," *Food Ethics* 1(4).

<sup>24</sup> In contrast, we might call the claim that meat eating is wrong on animal suffering grounds the *Weak Thesis*.

Finally, Section 2.5 wonders if we could save the person, but neutralise their prospective meat-eating harms. In this case, doing that would be the relevant lesser evil and we would be required to save them after all.<sup>25</sup>

As I noted above, while there are many twists and turns here, what the argument ultimately turns on is how bad is life in a factory farm compared to how good life is for a human, a topic on which people will differ substantially.

## 2.1. Comparing the Benefit to Someone of Living Longer against the Costs to Animals from a Meat-Eating Diet

Let's start by asking: Do meat eaters cause more unhappiness to animals through their diets than they experience happiness themselves in their own lives? By *happiness* I mean a net positive balance of pleasure over displeasure (unhappiness the reverse). Later on, I discuss whether valuing goods besides happiness might change the results.

One thing we want to know is how many years of animal life meat eaters cause to be created for each year they eat meat. Meat eaters might eat lots of animals, but to further simplify matters, let's assume meat eaters just eat chicken. Matheny and Chan point out that, of land animals reared for food, chickens constitute over 80% of all life-years.<sup>26</sup> They estimate that about 9 billion chickens were born each year in the US, which constitutes a total of 1440 million chicken life-years—chickens reared for meat live around 7 weeks, whereas egg-laying chickens live over a year. Given the human population of the USA, that means there are about 5 chicken life-years per human life-year.<sup>27</sup> While attempting to account for the other land and water animals eaten for food would presumably increase our tally of the negative effects on non-human animal lives, it would make matters far more complicated.

Given we have accepted, for the sake of argument, that meat eating is wrong, this means chickens' lives are, on average, bad. So, let's assume that the average chicken life is unhappy and the average human life is happy.

The next question to ask is this: What is the average level of human happiness compared with the average level of chicken unhappiness? This is a difficult question to answer. For our purposes, it's easiest to think about these as a ratio, which I write as *HW:AI* for short, that is, human well-being:animal ill-being. Using this 5-to-1 number, if the humans are less than 5 times happier than the chickens are *unhappy*, then, all else being equal, it is worse to save the life of an average meat eater than not save them. If *HW:AI* was 1:1, it would be 5 times worse to save the meat eater (all else equal).

---

<sup>25</sup> One thing I do not assess is the value of saving the life to the lifesaver. For instance, we can imagine someone claiming saving the life is the lesser evil, but only when they account for the unhappiness, guilt, despair, and perhaps the reputational costs, they would experience if they didn't save the person. I set these to the side: it is puzzling to think that I have a duty to save you, but *only* once when I realise that, if I don't, it will be bad *for me*.

<sup>26</sup> Matheny, G. and Chan, K. M. A. (2005) 'Human Diets and Animal Welfare: the Illogic of the Larder', *Journal of Agricultural and Environmental Ethics* 18(6): 579–94.

<sup>27</sup> I expect this ratio has not decreased since 2005. The USA's National Chicken Council (2021) estimates per-pound chicken consumption has gone up from 81.1 pound (2005) to 95.8 (2021), a nearly 20% increase; National Chicken Council. (2021). "Per Capita Consumption of Poultry and Livestock, 1965 to Forecast 2022, in pounds". Available at: <https://www.nationalchickencouncil.org/about-the-industry/statistics/per-capita-consumption-of-poultry-and-livestock-1965-to-estimated-2012-in-pounds> (accessed: 13 December 2021). The US population has grown from around 300 to 330 million over the same period.



We might attempt to answer it by asking ourselves whether humans' capacities for happiness are more than 5 times greater than those of chickens. We may doubt chickens can feel great happiness—they will never experience John Stuart Mill's 'higher pleasures' or Derek Parfit's 'the best things in life', whatever these are—and conclude meat eaters are at least 5 times happier than the average chicken is unhappy.<sup>28</sup> However, capacities are not necessarily relevant: we need to know the average *levels* of happiness and unhappiness that are experienced by humans and chickens, not what their *capacities* for happiness are.

To think about the levels, we could instead ask the following question, assuming we are as happy as the average meat eater: How happy am I during an average hour of my life, how unhappy do I think a factory-farmed chicken is during an average hour of its life, and what is the relative difference between the two? That may still seem hard to answer. Here are two alternative prompts which ask the same thing: Assuming how I feel *right now* is average, do I think I'm happier than an average factory-farmed chicken is unhappy right now, and if so, by how much? Alternatively: How many chickens, living in a factory farm for an hour, would it take to experience the same amount of unhappiness as I experience happiness in an average hour of my life?

While I sometimes experience great elation, most of my waking hours only feel mildly good. This is not because I think I'm suffering from some problem, either mental or physical, it is just that daily life doesn't come with very strong emotions. For reference, given how much time most of us spend working, how we feel during an average working moment is not far off the mean average happiness across our whole lives.

By contrast, the average moment a chicken spends in a factory farm believably feels *at least as bad* as I feel good now. From the footage I have seen of chickens in factory farms, that experience seems to be one of cacophonous chaos: animals packed into tiny spaces, pushing past each other to acquire food.<sup>29</sup> The descriptions of life as a broiler chicken are arresting:

Chickens have been bred to grow at grossly accelerated rates, causing a number of skeletal and cardiovascular problems. At the ends of their lives, they live at densities of around a square foot per bird, and 90% cannot walk properly, due to skeletal disorders.<sup>30</sup>

Quoting Singer's writing:

Chickens, reared in sheds that hold 20,000 birds, now are bred to grow so fast that most of them develop leg problems because their immature bones cannot bear the weight of their bodies. Professor John Webster of the University of Bristol's School of Veterinary Science said: "Broilers are the only livestock that are in chronic pain for the last 20 per cent of their lives. They don't move around, not because they are overstocked, but because it hurts their joints so much."<sup>31</sup>

---

<sup>28</sup> Mill, J. (1861). *Utilitarianism*. ch2; Parfit, D. (1986). "Overpopulation and the quality of life," in Singer, P. (ed.) *Applied Ethics*. Oxford University Press, p. 161.

<sup>29</sup> Mercy for Animals. (2011). "Farm to fridge—The truth behind meat production". Available at: <https://www.youtube.com/watch?v=THIODWTqx5E> (accessed: 18 June 2018); US Poultry. (2014). "Poultry insight: What is an AF0 and what is a CAFO?" YouTube. Available at: <https://www.youtube.com/watch?v=aq662XeMe3g> (accessed: 18 June 2018).

<sup>30</sup> Matheny and Chan (2005), p. 582.

<sup>31</sup> Singer, P. (2006). "Factory farming: A moral issue," *The Minnesota Daily*, 22 March. Available at: <https://www.utilitarian.net/singer/by/20060322.htm> (accessed: 18 June 2018). The quote from Professor

Regarding this last period of life, it is difficult to believe I experience more happiness than another creature experiencing chronic pain—I do not experience chronic pleasure.

Intuitions will differ, but it does not seem irrational, crazy for someone to conclude that the average human is as happy as the average factory-farmed chicken is unhappy—that is, that *HW:AI* is plausibly around 1:1. I am not sure how likely this is, but someone who thought it would not clearly be mistaken.

Given the earlier postulation that one year of human life creates around 5 years of chicken life, the interim conclusion is that meat eaters are causing around 5 times more unhappiness to animals than they are experiencing happiness themselves. Note, we have reached this conclusion not by claiming chickens have terrible, torturous lives and humans have wonderful, elated ones, but instead by thinking that much of human life is only mildly happy, whereas life in factory farms seems to be at best stressful, and for considerable periods, painful.

I consider three objections to this conclusion.

**Objection 1:** Quantifying average human happiness and comparing it to average chicken unhappiness is too nonsensically speculative even to attempt.

While tempting, this objection is not available to those who *already* hold the view that meat eating is wrong. Reaching that conclusion requires, at least implicitly, quantifying the value that individuals get from eating meat, comparing that to how much those animals suffer, and then concluding that imposing such suffering is impermissible. Hence, to hold we know enough about animal suffering to evaluate whether meat eating is wrong, but not enough to evaluate the Strong Thesis, is suspiciously *ad hoc*.

If we had a reliable means of making (cardinal) welfare comparisons across difference species, we would not have to engage in this speculation. However, I know of no such means, nor am I sure we could ever have one. We might conclude, on reflection, that we were mistaken to place any weight on our intuitive speculations and that we cannot sensibly compare welfare across species.

One issue though, if we thought that, is that we would then be pressed to abandon the Duty of Easy Rescue. Consider this plausible principle:

*Incomparability Thesis:* if two options are incomparable in value, both are permissible.

Suppose one accepts this. It, combined with the view meat eating is wrong, entails that we are permitted to either save or not save the stranger in the Shallow Pond. Yet the Duty of Easy Rescue holds we are required to do so. Inconsistency strikes.

The other issue is, of course, that if we can't make these comparisons, we now lack the justification to conclude meat eating is wrong in the first place.

**Objection 2:** We should be very uncertain about what *HW:AI* is.

Although I have listed this as an objection, it is not specific to this problem and has an uncontroversial solution: we should proceed as we do in other cases of uncertainty, by making a probabilistic estimate. We are quite happy to think that the average score of a fair, six-sided die is 3.5, even though we never roll a 3.5. What's more, if we are very uncertain, then we certainly cannot rule out the possibility that the Strong Thesis is true.

**Objection 3:** This analysis overestimates, for one or more reasons, the badness of the animals' lives relative to the goodness of the meat eaters'.

---

Webster is from *The Guardian*. (1991). "The meat factory: Cruel cost of cheap pork and poultry—factory methods have slashed meat prices in the last 30 years," 14 October.

There seem to be four different ways of pressing this, which I raise and address in turn. I call these *discounting* objections. My response is that, taken singly or collectively, it's unclear if they make enough difference to tip back the scales, assuming someone already believed *HW:AI* were around 1:1.

First, someone could simply claim that, when they reflect on it, *HW:AI* seems to be a lot larger than I supposed—humans are really a lot happier than chickens are unhappy.

I will not make a further comment on what *HW:AI* seems to be: intuitions will differ. Rather, I will say the following. For the objection to work, that is, to show that the meat eater experiences more happiness than they cause unhappiness via their diet, it cannot just be that *HW:AI* is a *bit* higher than 1:1 (assuming all else is equal). If the meat eater creates 5 chicken-years for each 1 of their own, *HW:AI* needs to be 5:1 or higher.

The second discounting objection is to point out that the only value I have accounted for is happiness, happiness is not the only thing with intrinsic value (although it may be one of them) and that including these other things changes the story. Put more technically, this objection rejects the conjunction of *hedonism* (the view that well-being consists in happiness) and *welfarism* (well-being is the only thing of intrinsic value).

There are three ways to press this objection, of which the strongest is to endorse an objective list theory of well-being (on which, for instance, goods such as friendship and wisdom are intrinsically valuable) and then argue that human lives contain relatively more of these goods than non-human lives. I note the two weaker alternatives momentarily.

We can pose the same query here as given a moment ago: Would moving to an objective list account alter matters enough?

Suppose for concreteness, well-being consists only in happiness, knowledge, friendship, autonomy and health. To make it better to save the meat eater, assuming *HW:AI* is 1:1, something like the following would have to be true: the average person's life contains some happiness, knowledge, friendship, love, achievement, and health and these all contribute to their well-being; the value the non-happiness components have together is four times larger than that of the happiness one. Remember, we are considering an average person here, not a genius scientist, someone with lots of friends, etc. Intuitively, the non-happiness components are not that important. To see this, consider the following case:

*Happiness or Other Goods:* You have two options: (A) 4 people have their happiness reduced to the 'neutral-point', where their overall level of happiness is zero; or (B) 1 person has their knowledge, friendship, autonomy and health reduced to their respective neutral points (whatever these are).

If we think the non-happiness items on the objective list contribute four times more well-being to someone's life than their happiness does, we would be indifferent between (A) and (B). However, it is not obvious that (A) is a better outcome than (B). This remains the case even if we add additional non-happiness items to the objective list.<sup>32</sup>

---

<sup>32</sup> I am grateful to an anonymous reviewer, who points out I am assuming a simple, additive account of well-being and there are other options. One alternative would be that non-hedonic goods only have value once a hedonic baseline is met but, if it is met, the non-hedonic goods comprise the vast majority of the value. I accept this is a non-trivial assumption but do not consider non-additive theories here: that would complicate matters enormously and the point would still remain that appealing to non-hedonic theories of well-being does not immediately allow someone to reject the Strong Thesis. That said, non-hedonists seem *more likely* to reject the Strong Thesis (which is unsurprising if we think humans instantiate more of the non-hedonist values).

The two alternatives are to appeal to a desire-fulfilment theory of well-being (on which your life goes better the more of your desires are met) or to take a non-welfarist stance (where certain things, such as beauty or knowledge are valuable intrinsically, independent of any effect on well-being).

Regarding the former, there are a variety of desire-fulfilment theories.<sup>33</sup> It is not obvious, on any of these, if they would particularly alter *HW:AI* compared to assessing well-being in terms of happiness. Is the extent to which humans are happy and factory-farmed animals unhappy radically different from the extent to which human desires are satisfied and factory-farmed animal desires are frustrated?

Regarding the latter, recall we are considering saving the life of a stranger—rather than (say) an eminent scientist—and it is unclear why we should expect someone picked at random to make much contribution these non-welfarist goods such as knowledge.

All this said, accepting one theory of value rather than another could change the relative importance of saving human lives to preventing unhappy animal lives.

A third sort of discount we might consider is one based on there being a morally significant difference between humans and non-humans, based on the kind of beings that they are. A crude and implausible version of this is a *pure species discount*, that is, a view on which it is more valuable to give the same well-being increase to humans than to non-human animals, justified purely based on species membership.<sup>34</sup> A more sophisticated and plausible alternative, one recently explored by Kagan,<sup>35</sup> would be to ground differential treatment on the basis of differential moral status, where animals with greater psychological capacities have greater status.

Again, we can make the same retort: if *HW:AI* is 1:1 then, unless the chickens' well-being is discounted by at least 80%, we would still conclude that it would be better not to save the human. I am unaware of any living philosophers who advocate such a large discount.<sup>36</sup> Shelly Kagan, who thinks such a discount *might* be justified, supposes it would only be small.<sup>37</sup>

A further problem for a steep species-related discount is that, however the discount is justified, it seems likely to generate unacceptable results in 'marginal cases'. Suppose the discount is based on the superior cognitive abilities of humans. Young children don't yet have these, which implies their suffering is unimportant or, at least, matters less. We might avoid this by saying that the capacity for rationality is what matters, rather than whether one has it at present. However, it now follows that humans with serious cognitive disorders—and who therefore lack these capacities—should have their suffering severely discounted, even though their abilities to experience pain, would be just as strong as those

---

<sup>33</sup> Heathwood, C. (2015). "Desire-fulfilment theory," in G. Fletcher (ed.), *The Routledge Handbook of Philosophy of Well-Being*. Routledge, pp. 151–63.

<sup>34</sup> Singer (2009).

<sup>35</sup> Kagan, S. (2019). *How To Count Animals, More or Less*. Oxford University Press.

<sup>36</sup> Kant, however, seemed to think we should give no weight to the interests of animals at all: "[the human being] is a being altogether different in rank and dignity from things, such as irrational animals, with which one may deal and dispose at one's discretion." Kant, I. (1798). *Anthropology, History, and Education*. Ed. R. Louden and G. Zoller. Cambridge: Cambridge University Press. Leaving aside how implausible this is, the target of this essay is not those—such as Kant—who would reject the Strong Thesis anyway.

<sup>37</sup> In print, Kagan claims a discount might—and also might not—be justifiable, but does not offer an account of how big it might be; Kagan, S. (2016). 'What's wrong with speciesism? Society of Applied Philosophy Annual Lecture 2015, *Journal of Applied Philosophy* 33(1): 1–21. At his 2016 Uehiro lectures at the University of Oxford, Kagan suggested the discount would only be very small.

of other humans. This is counter-intuitive. Finding a suitable rationale for this discount is not straightforward.

The fourth and final discount is based on the idea that meat eaters will soon switch to 'clean meat', grown from animal cells in a laboratory, or plant-based alternatives to meat, for example, burgers made of soy rather than beef. As no animals suffer in either process then, *prima facie*, the advent of clean and plant-based meat will end all the unhappiness that meat eating causes. Further, if the transition to clean and/or plant-based meat—the 'transition' for short—happens soon, the Strong Thesis will be false as meat eaters won't cause much more animal suffering.

Some readers will consider this discussion unnecessary. We could make this argument conditional on factory farming continuing and conclude it would cease to apply if and when that practice ends.<sup>38</sup> We could also note the peculiarity that, even if factory farming stops, it may nevertheless have been the case that, at earlier points, it would *have been* bad to save lives, even if it is no longer bad. However, it seems worth addressing this issue for readers who would otherwise dismiss the whole discussion as practically uninteresting on the grounds that the transition is around the corner.

One issue with appealing to this discount is that the transition is not certain to happen and even if it does it would not remove *all* meat eater-caused animal suffering. It might not be technologically possible to make meat that would be sufficiently cheap and/or tasty to prompt consumers to switch; presumably not all consumers would switch even if it were possible: one can, for example, already buy soy and rice milk, but most prefer cow milk.

The other issue is that, even if the transition removed all diet-related suffering, it is hard to believe it could happen soon enough. Suppose, for simplicity, the transition happens overnight. Let's say that Tim is the average stranger we're considering saving. Tim is the median age, 30, and might expect to live 40 more years. If *HW:AI* is 1:1 then, in 8 years, Tim's diet will have caused as much animal unhappiness as he will experience happiness for the rest of his life. Hence, the transition would need to happen within 8 years, which seems implausibly fast. I first started thinking about this topic in 2017. As I edit this final draft, it is late 2022. The transition seems no more imminent now than it did then.

Readers will recall that Shallow Pond supposes we're saving a child. Suppose this child is 10. For the same reasons as above, the transition would need to happen in 16 years to make it better to save the child. It is not obvious it will happen sooner than this, either. The situation may be one where saving children is good but—because the transition doesn't happen soon enough—saving adults over a certain age is expected to be bad.

We could use more realistic models for the transition, but such modelling exercises would be cumbersome and philosophically uninteresting without changing the basic point. Suppose the transition starts in *X* years and then 1/10 of all animal products are replaced each year over 10 years. Again, *X* needs to be soon to make saving lives better.

To summarise, taken individually, the discounts would need to be implausibly steep to make it better to save the meat eater, assuming someone already thought chickens were about as unhappy and humans are happy. However, someone might press further that, taken collectively, the discounting objections would be sufficient to change the result.

---

<sup>38</sup> I am grateful to this suggestion by Jeff McMahan.



I will make three quick replies to this. First, to state what should be obvious, we should be careful to avoid engaging in motivated reasoning and tweaking these so that we happen to get the intuitively palatable result; namely that saving lives is good.

Second, only some philosophers will be able to apply all four discounts, given their preferred theoretical machinery. For instance, Peter Singer has argued against a pure species discount and is (now) a classical utilitarian (that is, the right action is the one that maximises the sum total of happiness).<sup>39</sup> Therefore, Singer could only appeal to the first and fourth discounts, which makes the conclusion *relatively* harder to avoid for someone of his stripe.

Third, I fully accept that some will accept meat eating is wrong but, after giving careful consideration of the facts, nevertheless conclude the value of saving a hypothetical meat eater's life is greater than the negative effects of that person's diet. I leave the interested reader to crunch the numbers using their own inputs. My point is not that the Strong Anti-Carnism Thesis is true beyond reasonable doubt, only that is surprisingly plausible, opinions will differ sharply, and we could expect some to believe the thesis is true whilst others think it false.

## 2.2. Factoring in the Other Effects of Saving a Life

The effects a person's continued existence could have on the animals they eat is only one way in which saving their life could affect others. Three further effects are: (1) the grief caused to friends and family from a bereavement, (2) the effect on wider human society, (3) the effect on wild animals.<sup>40</sup> One might argue that, when we account for these impacts, it becomes more valuable to save lives—if they instead made it *less* valuable, that would support the Strong Anti-Carnism Thesis.

Effect (1) is likely to be small, relative to how bad it is for the person to die; I doubt many think the former is even 10% the size of the latter.

Effects (2) and (3) *could* be substantial considerations. However, it's neither clear that they are substantial nor whether they increase or reduce the value of saving a life. Investigating them is beyond the scope of this essay, so I don't attempt to include them. Arguably, this is quite a serious omission, and I will briefly elaborate on it.<sup>41</sup>

Regarding (2), many people seem to believe (human) overpopulation is a serious issue. However, as Greaves argues, it is not at all obvious, when we look at the matter closely, where we are in relation to optimum population—however we look at it.<sup>42</sup> The basic idea is this: it's reasonable to think there could be too many people: at certain point, we run out of scarce resources. Equally, there could be too few people: it is only when populations become large that they can support the specialised types of labour—for instance, physicians and metaphysicians—that make a high quality of life possible. Therefore, we can imagine there is some 'Goldilocks' global population size—not too big, not too small.

---

<sup>39</sup> Pertaining to discounts, see Singer (1975).

<sup>40</sup> This list would be exhaustive if it included effects on non-wild animals which are not eaten by humans, such as pets. I do not discuss these effects as they seem extremely small.

<sup>41</sup> I am grateful to anonymous reviewer for pressing me on whether we can justifiably consider the indirect dietary effects from saving lives but not all the other indirect effects.

<sup>42</sup> Greaves, H. (2022) "Optimum population size," in Gustaf Arrhenius et al. (eds.), *Oxford Handbook of Population Ethics*. Oxford Handbooks (online ed.). Oxford Academic. <https://doi.org/10.1093/oxfordhb/9780190907686.013.3> (accessed: 17 October 2022).

We then face the subsequent, serious empirical challenge of determining what the reality is. One possible avenue is to estimate the ‘carrying capacity’ of the Earth, that is, the maximum population size that the environment can support indefinitely, and assume the optimum population is around this. Yet the carrying capacity, at least for humans, depends on our technology: with better technology, we can feed a larger population, use less energy, and so on. Almost necessarily, we do not know what technology we will eventually develop. What’s more, according to the so-called ‘Boserupian account of innovation’, technological innovation is related to population size: as the population grows, we invent new technologies to deal with the problems caused by larger populations.<sup>43</sup> As a result, there is very little consensus among those who have attempted to determine carrying capacity. Greaves cites a survey of 65 such estimates in Cohen (1995), noting half of them are between 5 and 14 billion people, with a third above 20 billion.<sup>44</sup>

Even if we knew what the optimum population was, it would be a further matter to quantify, even vaguely, the value associated with adding or removing a life. As the sign and magnitude of the effect of saving life has on wider society are unclear, I set the issue to one side; it is outside the scope of the essay to make progress on.

It is similarly difficult to make confident predictions regarding (3) as it was for (2). The growth of the human population has reduced the proportion of the planet inhabited by wild animals. While most decry this as a tragedy, it is not clear if this is bad thing. Some, such as Tomasik claim life for wild animals is so tough that there is probably net suffering in nature and it would be better if there were less of it.<sup>45</sup> This is another factor that is too complicated and too uncertain to include here.

### 2.3. Adjusting for the Fact not All People Are Meat Eaters

So far, we have been wondering whether saving a randomly selected *meat eater* would be bad. The Strong Anti-Carnism Thesis concerns whether meat eaters have a sufficiently large negative effect via their diet that the value of saving the life of an *average* person is bad in expectation.

Adjusting the analysis to refer to a randomly selected person, rather than a meat eater, is unlikely to make much difference, at least in the developed world. However, 95% of people in America are meat eaters and 99% of animal products come from factory farms.<sup>46</sup>

---

<sup>43</sup> Boserup, E. (1976). “Environment, population, and technology in primitive societies,” *Population and Development Review* 2(1): 21–36.

<sup>44</sup> Cohen, J. (1995). *How Many People Can the Earth Support?* New York/London: Norton.

<sup>45</sup> Tomasik, B. (2015) “The importance of wild-animal suffering,” *Relations. Beyond Anthropocentrism* 3(2): 133–52.

<sup>46</sup> According to the National Chicken Council (2018) less than 1% of US chickens are free range. Sentience Institute (2018) estimate that 99.5% of land animals—meat chickens, egg chickens, turkeys, pigs and cows—in the US are reared in factory farms; National Chicken Council. (2018). “Chickopedia: What consumers need to know—The National Chicken Council”. Available at: <https://www.nationalchickencouncil.org/about-the-industry/chickopedia/#two> (accessed: 18 June 2018); Sentience Institute. (2018). “Sentience Institute US factory farming estimates—Google Sheets”. Available at: [https://docs.google.com/spreadsheets/d/1iUpRFOPmAE5IO4hO4PyS4MP\\_kHzkuM\\_-soqAyVNQcJc/edit#gid=0](https://docs.google.com/spreadsheets/d/1iUpRFOPmAE5IO4hO4PyS4MP_kHzkuM_-soqAyVNQcJc/edit#gid=0) (accessed: 18 June 2018). Note the vast majority of animals are chickens: Americans eat 8 billion chickens for meat each year, but only 90 million cows and 65 million pigs. A consumer report by GlobalData (2017) found 6% of American self-identify as vegans; GlobalData. (2017). “Top trends in prepared foods 2017: Exploring trends in meat, fish and seafood; pasta, noodles

Hence this, too, is a minor consideration and unlikely to be decisive for the Strong Thesis, at least in richer parts of the world where meat eating is common. The Strong Thesis would seem to be far less plausible in poorer parts of the world.

#### 2.4. What If Lesser/Greater Evils Are not Solely a Matter of the Consequences?

The preceding analysis assumed that, of two options, the lesser evil is simply the one which has the better consequences. For consequentialists, this is all that matters when deciding which evil to pick. However, some non-consequentialists will argue that other normative considerations matter here. In this subsection, I canvass the two options I expect non-consequentialists are most likely to appeal if they want to preserve the Duty of Easy Rescue. I argue that, even if one accepted these, it is unlikely they would alter the result.

The first consideration is that, when determining which option is the lesser evil, we must draw a distinction between the *direct* and *indirect* impacts of our actions and give more normative weight to the former. The distinction is that the direct impacts are those we personally cause to someone, whereas indirect impacts occur through the intervening agency of an agent who we directly impact. Consider:

*Surgeon:* We can either save 5 people with a drug or use that drug to save a surgeon who will save 100 different people.

47

Francis Kamm argues that “we should not ignore the 100 [whom the surgeon] could save. But we should also not give them their *full weight* as 100 individuals versus the 5 who need our resource to live.”<sup>48</sup>

The application here is that saving the stranger’s life is the direct impact, whereas the suffering they cause to the animal counts as an indirect impact, which gives it less weight when considering whether saving or not saving the person is the lesser evil. This makes it *relatively* easier to conclude we are required to save the stranger.

However, appealing to the direct/indirect distinction is unlikely to make much difference. Even Kamm seems to think we should discount indirect effects only slightly (“not give them their full weight”). Suppose we weigh indirect impacts 10% less. Hence, unless the harms that result from the average meat eater continuing to live are only *slightly* larger than the benefit they would gain, adjusting for directness won’t make saving the stranger into the lesser evil.

A second option would be to raise *anti-aggregationist* concerns, that is, the idea that many *small* goods(/bads) can never be morally equivalent to a *substantial* good(/bad). To

---

and rice; prepared meals; savory deli food; soup; and meat substitutes”. Available at: <https://www.reportbuyer.com/product/4959853/top-trends-in-prepared-foods-2017-exploring-trends-in-meat-fish-and-seafood-pasta-noodles-and-rice-prepared-meals-savory-deli-food-soup-and-meat-substitutes.html> (accessed: 18 June 2018). A poll conducted by Vegetarian Resource Group (2016) of 2,000 adult Americans found 3.4% ate a solely vegetarian diet, that is, they agreed with the statement “I never eat meat, fish, seafood, or poultry”; Vegetarian Resource Group. (2016.) “How many adults in the U.S. are vegetarian and vegan?” Available at: [https://www.vrg.org/nutshell/Polls/2016\\_adults\\_veg.htm](https://www.vrg.org/nutshell/Polls/2016_adults_veg.htm) (accessed: 18 June 2018).

<sup>47</sup> Kamm, F. (1998). *Morality, Mortality: Vol. I*. Oxford University Press.

<sup>48</sup> *Ibid.*, p. 108 (emphasis added).

illustrate, many believe that we should, for instance, save one life rather than prevent any number of headaches, even though the total suffering caused by the headaches would eventually be greater.<sup>49</sup>

The potential application of this approach to our case is to claim that preventing one human death is more morally important than preventing *any number* of chickens from existing and living miserable lives on the grounds that the harm to each chicken is so comparatively small. Hence, we are required to save lives in rescue cases as that is the lesser evil.

Even for those sympathetic to these anti-aggregationist responses, it does not appear at all obvious that the harm experienced over the life of factory-farmed chicken is small enough to be irrelevant. Why? We can illustrate this by comparison to the classic case of saving a life vs. preventing any number of headaches.

Suppose these hypothetical headaches last an average of one hour. Our factory-farmed chickens live, as noted, seven weeks. If they are awake 16 hours a day, that means they will have slightly over 750 conscious hours. It does not seem unreasonable to suppose that the average hour in a factory farm is *at least* as bad as the average hour of a headache. Admittedly, the chicken's level of consciousness is presumably lower at the start of their lives, but its experience progressively worsens as the animals grow and space reduces. Indeed, some readers will think an hour in a factory farm—sometimes several animals in each cage and cages stacked on top of each other, with the animals below barely able to move or to escape the defecation of those above—is more akin to an hour of torture than an hour of headache.

Therefore, even if we grant that no number of (average) headaches is as a great an evil as the loss of one (human) life, the case at hand is quite different: the harm experienced by the animal is several hundred times greater than that of a headache and an altogether different proposition. The principle that many small bads are not morally equivalent to a single large bad seems to have no application here.

To summarise this part of the discussion then, the two most obvious moves non-consequentialists might make to preserve the Duty of Easy Rescue, given the earlier concerns about the suffering caused by meat eating, will not take them very far.

## **2.5. Why not Save the Person and Neutralise Any Future Harms from Meat Eating?**

Someone might object that the above analysis was too narrow: it assumed the only choice is between saving or not saving the person. An alternative option, however, is to save the person and somehow neutralise the harms the person might later cause through their diet. One way to do this would be to rescue them, then insist that they stop being a meat eater. Another is to 'offset' their harms via donations to an animal welfare charity—you could make these donations for them or ask that they do it.<sup>50</sup>

Assuming this can be done at little cost—you have a quick word with them or donate a few pounds—it might seem ludicrous to let the person die, rather than save them and do this extra, small, neutralising act. Once we include the neutralising component, the Shallow Pond case is now, more or less, back to how we had originally conceived of it: you can prevent a large harm by saving a life and incurring a small cost for yourself;

---

<sup>49</sup> Norcross, A. (1997). Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs*, 26(2), 135–67

<sup>50</sup> I am grateful to Jeff McMahan for raising the former possibility and to an anonymous referee for the latter.

therefore, morality requires that you save the life and bear those costs. Clearly, the least evil is to rescue the person and neutralise the harms, so you are required to do that.

This response might seem like a neat fix, one that makes the concerns raised in this essay disappear. However, it is worth pointing out that, if this is the correct response, just how peculiar and counter-intuitive a revision of morality it would be. If we are really required to rescue people *and* then neutralise their harms, it means we are not permitted *just* to save the person: you can't simply drag the person out of the pond and get on with your life; you must do something else, such as lecture them about the horrors of factory farming or write a cheque to an organisation like the Humane League.<sup>51</sup>

Furthermore, there are serious problems with both neutralising options.

With regard to the 'rescue and convert to veganism' option, whether this would result in a lesser evil than not saving the person depends on your chances of success. Of course, if you knew you would persuade them, then 'rescue and convert' would be the lesser evil and thus possibly be required. However, it seems most likely they would assume your request was mad and ignore it ("You won't believe what happened to me today. I fell in the pond and would have drowned if someone hadn't pulled me out. But that wasn't the weird thing. The person who pulled me out then asked if I ate a lot of chicken and demanded I stop.") The reason we seriously countenance not saving the Drowning Dictator is that, while the best outcome would be if you saved him and then successfully convinced him to stop doing bad things, we recognise this outcome is not at all likely.

There is *some* chance that rescuing and requesting a diet change would work, so we can conceive of it as another sort of 'discounting objection': as a simplistic example, if you expect that if you ask someone to stop eating meat, there's a 10% chance they will give up entirely and a 90% chance nothing will change, then you would reduce the expected effects of their diet by 10%. Therefore, given someone's other beliefs, that could be enough to tip the balance into making 'save and convert' a lesser evil than not saving the person.

Let's turn to offsetting. Many believe it is permissible to act in ways that would otherwise be wrong if we *offset* the act with another of comparable magnitude. There are various types of offset,<sup>52</sup> and the most plausibly permissible is that where the offset prevents a harm from ever occurring. For instance, if I fly from London to New York and plant enough trees in my backyard to absorb the equivalent in greenhouse gases, then, quite possibly, no one is ever harmed, so my flight is permissible—whereas it wouldn't be if I hadn't offset the greenhouse gases. This case is much like one where I drop an iron that would have landed on your foot, but I catch it before it gets there—no harm done.

However, as John, Askill, and Wilkinson point out, nonconsequential moral theories generally regard it as impermissible to harm one group, but then proportionately benefit another: I am not permitted to torture a few people, but then offset that by (say) donating enough money to Amnesty International's anti-torture campaigns to prevent an equivalent amount of torturing: helping other innocents doesn't make my torturing permissible.<sup>53</sup> In

---

<sup>51</sup> We might alternatively think we are under a conditional obligation: you are not required to save the person, but if you do, you must then act to neutral their harms. I don't consider this possibility here. For one, it abandons the Duty of Easy Rescue. For another, it's puzzling to think how we can have such obligations. For some discussion, see Pummer, T. and Muñoz, D. (2022). "Supererogation and conditional obligation," *Philosophical Studies* 179: 1429–43.

<sup>52</sup> For a taxonomy and discussion, see John, T., Askill, A. and Wilkinson, H. *The Offsetting Paradox*. Unpublished manuscript.

<sup>53</sup> *Ibid.*



this case, donations to some animal charity, even one that focuses on factory farming, will inevitably affect the lives of different animals from the ones the meat eater would consume. Hence, John, Askill and Wilkinson argue non-consequential moral theories face a dilemma: they must either allow any act, no how matter how grievous, to be offset or accept an implausibly sharp discontinuity between offsettable and non-offsettable actions.<sup>54</sup>

Offsetting is not permitted on consequentialist theories. Suppose you can do (A), an act with bad consequences, and, quite independently, you can offset those harms by doing (B), an act with very good consequences. Doing both (A) and (B) is overall better than doing neither. However, because not doing (A) and doing (B) leads to the best outcome, that is what consequentialism demands. To illustrate, suppose donating to Amnesty International is the best use of my money. Clearly, I ought to give to Amnesty International and not torture people, rather than give to Amnesty *and* torture people. Returning to our case, rescuing the person and donating to animal charities are separate acts. If we think not saving the person is better than saving them and that donating to animal charities is the best use of your money, then quite clearly consequentialism requires that you should both not save the person and donate the money.

One possible, if flimsy, response to this would be for someone to claim that they will donate to the animal charity *if and only if* they save the person—they might say that if they save the person, then they will think about the additional animal suffering that person may now cause, and that will motivate them to donate something. Therefore, they claim that, at least for them, saving and offsetting really is the best outcome. However, that any such people exist, and that they couldn't motivate themselves to give by use of other, less harmful methods—such as, for instance, by thinking about (animal) suffering—seems unlikely.

### 3. If the Strong Thesis Is True, Must We Abandon the Duty of Easy Rescue?

The previous section asked whether the Strong Thesis is true and argued there is a case for it, one that some will believe, but that it is by no means certain. In this section, we will assume that the Strong Thesis is true and ask if we must therefore give up the Duty of Easy Rescue.

The problem, recall, for the Duty of Easy Rescue is that, if one accepts the Strong Anti-Carnism Thesis and that there are lesser-evil justifications for doing(/allowing) harm, then we are not required to save lives in rescue cases. This is inconsistent with the Duty of Easy Rescue, which holds that we are. To avoid inconsistency, we must give up on something. We have now discussed the Strong Anti-Carnism Thesis at length. I take it that the existence of a lesser-evil justification for doing(/allowing) harm is not up for debate. That pushes us to abandon the Duty of Easy Rescue.

That we are not required to save lives when we can do so easily may seem highly counter-intuitive, even repugnant. Let's now revisit the likely origin of this intuition and remind ourselves of the argument; after this, jettisoning the Duty of Easy Rescue may not seem so unintuitive.

The Duty of Easy Rescue does not seem to be a fundamental moral principle. Rather, it derives its intuitive force, at least in Shallow Pond, from two underlying premises: (1)

---

<sup>54</sup> Ibid.

saving a life is overall very good, and (2) we have a duty to promote the good when we can do so easily. Saliiently, (1) is a mixed empirical and moral claim whilst (2) is a purely moral claim. What explains why we do not think we have to save the life in Drowning Dictator is that the empirical facts are different and (1) is not true: we expect that saving the dictator's life would lead to much suffering, and because suffering is bad, saving the life is bad overall. It seemed obvious, prior to accounting for humans' effects on animals through their diets, that rescuing the person in Shallow Pond was very good. However, if we really did believe the Strong Anti-Carnism Thesis were true then, once we account for this, *not* saving lives in ordinary rescue cases is now the lesser evil and so the Duty of Easy Rescue is robbed of its original plausibility.

How should we revise the Duty of Easy Rescue? Given that the worry is about animal suffering, and assuming the Strong Thesis, it seems it should be *something* like:

*The Revised Duty of Easy Rescue:* In rescue cases, we are not required to save lives. However, in the unlikely event that we are confident the person is not a meat eater (for instance, they are vegetarian or vegan), we are required to save them.

While we may find this conclusion distasteful—I certainly do—that alone does not give us any reason to reject it.

#### 4. Potential Additional Implications

The unexpected and counter-intuitive conclusion that the impact of factory farming could be so large raises several questions—here, I consider four. These could all be discussed at substantial length but for reasons of space these can only receive brief treatment.

The initial three concern what may follow if we now conclude that the existence of humans is overall bad. The fourth relates to the implications for effective altruists.

First, if the Strong Thesis is true, would we be required *not* to save lives in rescue cases—rather than merely not required to do so?

If the Strong Thesis is true, not saving strangers' lives in rescue cases is the lesser evil. Thus, for consequentialists, as for non-consequentialists such as Frowe who hold that we are required to do the lesser evil—at least when it is not too demanding—it follows that we are required *not* to rescue strangers.<sup>55</sup> Peter Singer, a consequentialist, would therefore be pushed to say it is wrong to save the drowning child in the Shallow Pond if he accepted the Strong Thesis (which, of course, he might not).

Second, if saving lives is overall bad, are we permitted, and perhaps even required, to kill other people to prevent their meat eating?

This does not straightforwardly follow. Non-consequentialists may appeal to the familiar doing—allowing distinction, as noted earlier: we are not permitted to cause one harm unless the harm we prevent is substantially greater.

What's more, consequentialists and non-consequentialists alike will point out that, in general, going around killing people is exceedingly unlikely to be how one does the most good—*inter alia*, you can do little good from inside prison. For those concerned with factory-farmed animal suffering, presumably the better route is to raise that issue publicly and support campaigns against it—not to kill a handful of people.

Third, even if it's wrong to kill people, should we be trying to reduce the overall human population without killing people? Perhaps we could try to shut down hospitals, so lives

---

<sup>55</sup> Frowe (2018).

are not saved, or aim to make birth control more widespread, so fewer new people are born.

The answer here is much the same as the one above. Even if it would be better, in general, if there were fewer people, it would seem we have only a *pro tanto* duty to realise this, one which would be easily overridden by other considerations, most saliently that there is something even better we could do instead.

Fourth and finally, we can ask: How might accounting for humans' practices on animals alter our priorities if we are trying to be *effective altruists*, that is, aiming to do the most good with our spare resources? Two options commonly recommended by members of the effective altruism social movement are to give to organisations in the developing world that save lives or reduce poverty.<sup>56</sup> However, the longer people live and the richer they are or become, the more meat they eat and the more suffering they will cause as a result. Studies show very consistently that, as people get wealthier, they consume more meat.<sup>57</sup> Hence accounting for the effects that human practices have on animals will reduce the value of both of these by *some amount*—it won't necessarily make such actions negative, not least when we consider that those in the developing world eat much less meat. Of course, opinions will differ over how large this reduction is, something that will vary, as noted, between countries. Crucially, this reduction does not rely on the truth of the Strong Thesis: just so long as we think humans have *some* negative impact on others, it applies.

## 5. Conclusion

I have argued that two widely accepted beliefs—we ought to save strangers and we ought not to eat meat from factory farms—are in a deep and underappreciated tension. While we would not normally consider these beliefs to be relevant to each other, I pressed the straightforward problem that, if we have those animal welfare concerns then, when we account for them, it reduces, and may remove, the obligation to rescue others.

I consider this surprising and disturbing. It brings to mind the words of an Italian saint, popularised by British Prime Minister, which, suitably rephrased, aptly captures the occasional role of philosophers:

*Where there is harmony, may we bring discord.*

*Where there is error, may we bring truth.*

*Where there is faith, may we bring doubt.*

*And where there is hope, may we bring despair*

**Acknowledgments:** I am grateful to the following for their help with various stages of this document: Joe Bowen, Tommy Francis, William Gildea, Caspar Kaiser, Patrick Kaczmarek, Todd Karhu, James Kirkpatrick, Kacper Kowalczyk, Michael Masny, Aiden Penn, Theron Pummer, Gary O'Brien, William Rooney, Korbinian Ruger, Julian Savulescu, Peter Singer, and Rhys Southan. I thank Doris Schneeberger for raising with me the question of whether factory farming might cause so much suffering that saving the average human life would be bad.

---

<sup>56</sup> MacAskill, W. (2015). *Doing Good Better*. Faber & Faber; Singer (2015).

<sup>57</sup> Delgado, C. L. (2003). "Rising Consumption of meat and milk in developing countries has created a new food revolution," *The Journal of Nutrition* 133(11): 3907S–3910S.