

The Diner's Defence:

Producers, Consumers, and the Benefits of Existence

Abelard Podgorski

Abstract: One popular defence of moral omnivorism appeals to facts about the indirectness of the diner's causal relationship to the suffering of farmed animals. Another appeals to the claim that farmed animals would not exist but for our farming practices. The import of these claims, I argue, has been misunderstood, and the standard arguments grounded in them fail. In this paper, I develop a better argument in defence of eating meat which combines resources from both of these strategies, together with principles of population ethics, and discuss its implications for which sorts of meat it is permissible to eat. According to the *diner's defence*, there is an asymmetry between producers and consumers of meat. Producers can prevent the suffering of animals without preventing their existence, but consumers cannot. This asymmetry grounds a defence against harm-based objections to eating meat which is available to the consumer alone, and which avoids the controversial commitments about moral status or the interests of nonhuman animals endemic to existing attempts to justify omnivorism.

Keywords: Vegetarianism, Animal Ethics, Population Ethics, Non-Identity

1. Introduction

In this paper, I will present a qualified defence of the consumption of animal meat.¹ Most existing defences of omnivorism rely on controversial claims about moral status – claims that I believe have been convincingly shown to be deeply contrary to moral common sense.² Some rely on views about the interests animals have in their own future that are questionable on philosophical or empirical grounds.³ Some depend on accepting one or other specific comprehensive moral doctrine.⁴ The well-worn disputes over these issues has, in my judgment, reached an impasse, or in some cases a position favourable to the vegetarian, and my aim here is to give a defence that bypasses them entirely.

The thesis will be in some respects modest. First, the challenge is aimed at just one important kind of reason often adduced against consuming meat – what we might call a *harm-based* reason –

¹ I will focus on the consumption of animals for food, but the argument will generalize in straightforward ways to the use of other animal-based products, like leather, and to the use of products that are not made from animals but which use animals in the process of their production.

² Norcross [2004] provides a particularly clear presentation of the vegetarian's case.

³ See Harman [2011] for a discussion of philosophical worries about this strategy.

⁴ For example, Narveson's [1987] contractarian defense.

grounded in the contribution eating meat makes to the harms that are inflicted on future animals.⁵ Appeal to this sort of reason is the most common strategy among opponents of eating meat, but not the only one. I'll briefly consider other reasons why eating meat might be wrong at the end of the paper.

Second, it will only be a defence of the *consumption* of meat,⁶ and not a defence of the farming practices themselves. It is for this reason that I will call this the *diner's* defence. It will be compatible with the diner's defence that *no* possible farming practices are morally permissible. These distinct moral questions are sometimes conflated, or it is assumed that the harm-based reasons against the production of meat carry over, in some straightforward way, to its consumption. I will try and show that the gap is not so easy to cross as it may seem.

Finally, I will only be defending consuming *some* types of meat – in particular, meat produced under conditions where the animals have lives that are worth living. This should not, however, be confused with a defence only of consuming meat that has been produced *humanely*, that is, under morally acceptable conditions. The requirement that animals have lives that are worth living is a lower bar (though, tragically, almost certainly higher than most existing farming practices manage to clear). Morally horrific practices may still produce lives that are on balance worth living. Which animals clear this bar will be a distinct and difficult question – I will make some suggestions towards the end of the paper, but the reader need not agree with them to appreciate the main point.

While modest in these respects, my argument will be unusual in making the most extreme concessions possible to the vegetarian on essentially all the main points of contention in the literature. I will grant that animals have moral status – and indeed, that they have *exactly the same* moral status as human beings, including rights of the same kind and importance. I will grant that animals have an interest in their future, and that they are harmed by their deaths – indeed, that their deaths are *just as bad for them* as our deaths are for us. I will grant that the act of purchasing meat causes, or substantially risks causing, animals to be seriously harmed, and that the benefits we get from eating meat are

⁵ These harms may be problematic for consequentialist or for deontological reasons – this paper means to remain neutral about the deeper explanation for the wrongness of harm.

⁶ By consumption I include both the purchase and eating of meat. Though the discussion is often framed in terms of decisions about what to eat, as will become clear, it is the purchase that is the most plausible source of harm.

incomparably less morally significant than the harms inflicted on farmed animals. I will grant that whatever the harms caused by alternative eating practices (for example to field animals who are killed in the harvesting of plants), they are incomparably less morally significant than the harms inflicted on farmed animals. I will not claim that animals are *replaceable* – that one can offset the harm of killing an animal by replacing it with an equally happy animal afterwards. Finally, I will make no empirical assumptions about the psychological capacities of animals. No existing defence of omnivorism, as far as I am aware, hands over this much ammunition without a fight.

My argument will be grounded in two observations that occasionally arise in debates over vegetarianism. First, the fact that the causal relationship between purchasing and eating meat and the suffering of animals is indirect – the harms that animals suffer are not directly inflicted by those who consume meat but rather by the farmers who raise and slaughter the animals (call this *Indirectness*). Second, the fact that animals which are factory farmed would not exist but for those practices (call this *Existence-dependence*). The most natural arguments built from these premises and defended in the literature, however, are deeply flawed, for reasons I will explain. I believe these facts *are* morally relevant. But they only matter *together*. Consumers and producers of animal meat stand in different causal relationships to the existence and suffering of animals, and, I will argue, this has crucial moral significance. Producers of animal meat can prevent the suffering they inflict on animals without preventing their existence. Consumers cannot.

2. Indirectness

Before we begin the main argument, which depends on both *Indirectness* and *Existence-dependence*, it will be helpful to look at each in turn, and at the typical way these claims are conscripted in defence of omnivorism. There is a lesson in the failure of each of these arguments which will help illuminate the path towards the stronger defence of this paper.

First, let us consider the fact that the consumer's relationship to the death and suffering of animals in factory farms is indirect. A simple line of reasoning goes as follows. Grant that the way

animals are treated on factory farms is morally wrong. Nevertheless, the people who consume meat do not *themselves* slaughter any animals. They never pushed any animals into cages too small to accommodate them. They did not choose the equipment used to produce the meat they consume. The blood of the animals, the argument goes, is on the hands of the farmers alone.

If one teaches long enough, inevitably one will run into students who make this kind of argument. But no self-respecting philosopher should accept it, and as far as I am aware no self-respecting philosopher does. The problem is not difficult to spot – on no plausible view are we morally responsible only for the harms we *directly* inflict. If I pay somebody else to commit murder, the fact that I was not the one who held the knife does not let me off the hook. Likewise, I could not escape criticism merely by paying someone else to see to it that animals suffer for my benefit.

So, the mere fact of indirectness does not excuse us from responsibility for the unnecessary suffering of animals on factory farms. But a more sophisticated omnivore might try to do better. The indirectness of the diner's relationship to the animals they eat, one might claim, means they do not cause the animals' suffering *at all*. The animals that they purchase are already dead, after all. It is too late for anything they do to affect those animals' fate. And the individual diner is so far removed from the decisions that affect the production of animals that they cannot expect their choice to have any effect on the suffering of future animals either. We can call this the causal impotence argument, and unlike the naïve argument above, some philosophers have defended versions of it.⁷

But again, the vegetarians have a response.⁸ When you make a purchase, the standard reply goes, you signal demand for the product. The producers respond to this demand in aggregate by increasing production. For straightforward economic reasons, the amount of meat produced for consumption will roughly track the amount consumed. Now, it is true that producers do not respond to signals on an animal-by-animal basis. When a certain threshold of demand is met, they will increase production by a substantial amount, and it is impossible to predict exactly where this threshold is, or whether one's individual purchase will cause the threshold to be met. But while this may mean that any individual

⁷ For instance, Harris and Galvin [2012].

⁸ See Singer [1980] and Norcross [2004].

purchase is very likely to have no effect,⁹ this is offset by the fact that in the case where it does have an effect, this effect may be very large. If one in every ten thousand chickens purchased causes demand to hit a threshold that triggers the farming of another ten thousand chickens, then while it is very likely any individual chicken purchase does nothing, the *expected* contribution, calculated by multiplying the chance your action has an effect by the magnitude of the effect, still amounts to approximately one additional farmed chicken per chicken purchased.¹⁰

While some of the defenders of the causal impotence argument have tried to raise challenges for this response,¹¹ I find it compelling. Undoubtedly, the indirectness of our relationship to the suffering of animals is *psychologically* significant. It is difficult to imagine that people would be willing to eat meat produced under modern farming conditions if they had to inflict all of the harms directly. But indirectness by itself does not acquit us from blame, and the nature of the indirect relationship is not one that renders our causal contribution – or at least our *expected* causal contribution – insignificant.

I take this dialectic to be won by the vegetarian. For our purposes, however, the main lesson to learn from this is that the contribution to the harms inflicted upon animals one makes by eating meat is that one increases demand, causing (or risking causing) more animals to be produced, to experience suffering, and to subsequently be slaughtered. This will become important later.

3. Existence-dependence

Now let us set the preceding discussion aside for the time being, and consider a second, independent argument in favour of omnivorism, which, unlike the argument from indirectness, purports to justify the entire practice of farming and eating meat. According to this line of reasoning, farming animals for food is permissible because those animals would not exist but for our farming

⁹ Even this is not obvious; as Norcross [2004] notes, even if it is not itself a threshold purchase, it may cause some future threshold to be met earlier than it otherwise would.

¹⁰ Matters are slightly complicated by the fact that higher consumption raises prices, offsetting to some degree the effect on production of eating more meat. But this effect isn't large enough to substantially affect the argument. Incorporating information on price elasticity, Norwood and Lusk [2011, pg. 223] estimate that, on average, an extra pound of meat eaten increases production by 0.68-0.76 pounds).

¹¹ See e.g. Budolfson [2018].

practices, and they are benefited by existence. As Leslie Stephen puts it, writing over a century ago:

Of all the arguments for Vegetarianism, none is so weak as the argument from humanity. The pig has a stronger interest than anyone in the demand for bacon. If all the world were Jewish, there would be no pigs at all. [1896]

More recently, a defence along these lines has been endorsed by Loren Lomasky [2013].¹² Notably, this sort of defence could only apply to certain farming practices – those in which the lives of animals were at least worth living. If an animal experienced nothing but agony throughout its short life, it would not be plausible to claim that the animal is benefited by existence. Stephen, who was writing before the advent of modern factory farming practices, could perhaps be forgiven for not making this explicit.

There is a way of understanding the existence-dependence argument which is *impersonal* and a way of understanding it where it is *person-affecting* (or, if we don't want to sound like we're making assumptions about the personhood of animals, *animal-affecting*).¹³ On the impersonal interpretation, farming practices are justified because the existence of animals in those practices makes *the world* a better place. On the person/animal-affecting interpretation, farming practices are justified because the existence of animals in those practices is good *for those animals*.

According to a view like classical utilitarianism, permissibility is determined entirely on impersonal grounds. Peter Singer, though generally opposed to the farming of animals, suggests that such a view could justify a limited omnivorism [1993, Ch. 5]. On Singer's view, the existence of happy animals is an impersonally good thing – any disvalue in killing them comes from the loss of additional impersonal goodness their continued lives would contribute. Consequently, whatever disvalue there is in their death can be offset as long as they are replaced by new, happy animals.

I'm going to set aside this impersonal version of the argument, however. Although some

¹² McMahan [2008] considers an argument of a similar sort, though he does not ultimately endorse it.

¹³ I borrow this distinction from Parfit [1984].

comprehensive moral views license this sort of reasoning, in this paper I am concerned with arguments that do not assume any particular theory. And even within a broadly utilitarian framework, it is controversial both whether the mere existence of more happy animals makes the world impersonally better, and whether harms inflicted on one being can be offset by creating other beings with good lives.¹⁴

On person-affecting versions of the argument, which will be the focus of this paper, it is not assumed that harms inflicted upon one animal can be offset by benefits to a newly created animal. It is important to these arguments that the animals *themselves* are benefited by the practice, and not just that the world as a whole is improved. If the animals have lives worth living, and their existence depends on farming practices, the thought goes, then the practice benefits them.¹⁵ And if they are benefited by the practice, they do not have grounds for any *moral complaint* against it. We benefit from the practice, since we get to enjoy eating meat. The animals benefit, since they get to exist. So, it's a win-win.

I'll argue later that there is something to the existence-dependence idea. But it cannot be used to justify our practices in the way its proponents suggest. I'll start by pointing out that it generalizes in absurd ways. Then I will give a diagnosis of the mistake.

Imagine that you discover a culture with the following unusual practice. Each set of parents treats their first two children very well, with all the love and affection that parents give to their children in our world. These children live full, happy lives, and carry on the family name. But after they have secured their bloodline, most families continue to have more children. The parents have these children only because they enjoy abusing them. The extra children live in squalor and are regularly beaten violently by their parents and older siblings. When they become old enough to fight back, they are killed. Nevertheless, the beatings are infrequent enough that the children's lives are worth living, and

¹⁴ The canonical presentation of these issues in population ethics can be found in Parfit [1984]. See Visak [2013] for a discussion specifically with respect to animals and within the context of utilitarianism.

¹⁵ One might worry about the claim that beings *benefit* from having a good existence, on the grounds that had they not existed, they could not have been *worse off*, since they simply would not have any quality of life at all [Bramble 2015]. But as McMahan [2008] and Parfit [2017] point out, it is plausible that existing with a good life can be good for a person in a *noncomparative* way. So we should allow that there can be existential benefits (to people who have lives that are worth living) and existential harms (to people with lives that are not worth living). In any case, the reasoning that denies that animals would benefit from existence would also deny that animals can be harmed by existence, and this would be enough for the person-affecting strategy to challenge the existence of harm-based reasons against omnivorism.

they are, overall, glad to be alive.

If the existence-dependence argument above is good, then it seems this culture's treatment of their third and fourth children is justified. After all, if the practice of abusing children did not exist, the parents would stop at two children and have no more, and none of the abused children would exist at all. Their lives are worth living, so they are benefited by existence. By parallel reasoning, then, they have no complaint against their treatment, and the harms they suffer are not a moral strike against it. This consequence, I take it, is absurd.

In response to a worry that his argument would overgeneralize to human beings, Lomasky appeals to a familiar line about the different status of humans and animals with respect to rights. While the practices above do benefit the humans under them, he suggests, they violate rights that humans, but not animals, possess – for example, the right not to be killed. It is not permissible to violate these ‘rightful liberties’ even for a person’s own good.

I do not think that this is a very satisfying response. First, the view that *no* benefit to a person could override the prima facie right not to be killed is contentious, and here the benefit is of the most substantial sort - it is *having a life at all*. Indeed, part of what the people under these practices gain by existing are the very conditions for having rights – the rationality, for example, which on many such views gives human beings their special value. So, it is far from obvious that a similar justification would not be available even if we take rights on board.

More importantly for our purposes, however, this strategy would make the success of the existence-dependence argument parasitic on controversial claims about rights and moral status, which is something the omnivore would prefer to avoid. The position of the argument is a bit precarious. On the one hand, the argument is only necessary if animals have a moral status that requires us to respect their interests. On the other hand, it would only avoid the objection if animals have a substantially different moral status than human beings - different, moreover, in precisely the right respects to block the generalization of the argument.

In any case, I think we can refute the argument more directly. The problem with the existence-dependence argument as stated is that to justify on harm-based person-affecting grounds a practice that

involves the creation of people or animals, it is not enough to show that the existence of *the entire practice* is better for those under it than the nonexistence of the entire practice. As Bramble [2015] points out, farming practices involve multiple parts – bringing animals into existence, and then treating them in a particular way. The practice is not justified unless both of these parts are justified. And the second part can be wrong, in virtue of giving worse lives to the animals than they might have had [pg. 138] even if the first is a benefit to those animals.

Let Theodore be one of the third-born in our thought experiment. It is better for Theodore that the entire practice has existed than that none of it exist. But take any particular act of abusing Theodore. Theodore does not benefit from that act. And his existence does not depend on *that* act being performed. Perhaps Theodore would not have been born were it not for the earlier existence of the practice and the abuse of other children that entailed, which gave his parents the intention to have a child to abuse for themselves. But that does not make it okay, as far as Theodore's interests are concerned *now*, once he already exists, to continue the practice, and carry out the rest of the parents' harmful intentions. If someone's parents would never have met but for being slaves, this takes no weight away whatsoever from their objections to being enslaved themselves.

Similarly, it is perhaps true that the pigs being killed or abused on farms today would not exist but for the history of that practice. But the most this shows is that they do not have a complaint, on their own behalf, against the *earlier* acts in that practice. Now they have already been born. The farmers are not compelled by history to continue the practices or carry out the intentions that motivated them to cause the pigs' birth. And the pigs do have a complaint against their current treatment, because the farmers had options available to them which would have led to less suffering and/or a less untimely death for those pigs.

4. The Diner's Defence

Let us take stock of what we've learned so far. We saw that the mere fact that our causal relationship to animal suffering is indirect does not by itself offer absolution for the suffering of

animals, and that our actions are not causally impotent with respect to the suffering of animals. Our actions create demand which leads to animals being bred, farmed, and slaughtered, in a way that is significant even once we take the chances of making a difference into account.

We also saw that the fact that the animals being farmed would not exist but for those practices could not justify the conditions of their farming, because while the animals may not have a moral complaint against the acts involved in the practice which were necessary for their existence, they can have complaints against the acts that are unnecessary for their existence, including all of the choices made after their birth about their treatment.

Now I want to propose a way to draw together these lessons to produce a successful defence of eating meat. There is an insight in the existence-dependence argument that I think is correct. It is this: *an animal or person does not have a harm-based complaint against an act which causes them to exist, if no alternative act would have led to their existence, and their existence is good for them.* This is true even if their existence involves many wrongful harms being inflicted upon them. Children who live good but imperfect lives do not have a complaint against their parents for bringing them into existence. In a sense, the parents are causally responsible for all the suffering in their children's lives – there would be none if the parents had not conceived. But likewise, they are responsible for all the happiness in those lives. The children might have something to complain about if the suffering were not outweighed, but otherwise, they have no complaint arising from their own interest against the acts that brought them into being.

The children may, of course, have other legitimate complaints about their treatment by their parents. If their parents abuse them, like the third-born in our earlier example, the children have a complaint against that. And this is why existence-dependence cannot be used to justify the entire practice, in both Theodore's example and in the case of factory farming. For not all acts involved in those practices satisfy the existence-dependence condition, properly formulated above.

Here is where indirectness steps in. The nature of the indirect relationship between the consumer and the treatment of animals implies an asymmetry between producers and consumers. This is the second crucial point. *Producers of animal meat have control over whether animals suffer and die which comes apart from their control over whether those animals exist. Consumers of animal meat do not.*

This second claim rests on an empirical assumption - that the effect of the diner on demand is largely realized through the creation of new animals which are farmed, and not through the increased mistreatment of individual animals who already exist or will exist no matter what the diner does. For some possible practices, such as those in which meat is provided through the hunting of wild animals, this assumption may not be true, and the diner's defence will not apply. But I think there are good reasons to suspect that in modern practices, this assumption is true. The lives of factory farmed animals are heavily scheduled, and it is unlikely that increased demand would mean any significant number of animals destined for better things are redirected onto someone's plate, rather than new animals being bred. Moreover, it is implausible that any *existing* animals will be affected, because of the time it takes for the economic signals sent by the diner to serve as inputs into a decision and for that decision to be implemented, compared to the relatively short lifespans of farmed animals. Finally, unless the conditions of animals' creation are causally isolated from decisions about the conditions in which they live, any major overhaul of farming conditions potentially brought about by a diner's choices will be enjoyed almost exclusively by animals that would not otherwise exist. I think this is enough to justify a presumption that existing defences of vegetarianism have not even attempted to undermine.

To see concretely how the argument works, take an individual pig, Dickey, who is born, raised, and slaughtered. Dickey's existence, his suffering, and his death are the causal product of both the actions of a farmer (or many farmers), and the actions of a diner (whichever diner, we might say, sent with their pork purchase the threshold demand signal which led the farmer to increase production, leading to his birth). With respect to the farmer, Dickey's existence is the product of one set of choices - for instance, the decision to breed his mother. His suffering and death are the product of another - the decision to use small cages, the decision to slaughter him, and so on. With respect to the diner, however, there is just one decision - the decision to purchase pork, which is causally responsible both for Dickey's existence, and for his suffering.

What this means is that there is something the farmer could have done which would have prevented Dickey's suffering and untimely death without preventing his existence. But there was

nothing that the diner could have done which would have prevented Dickey's suffering or untimely death without preventing his existence entirely. This is morally important. Dickey has a complaint against the farmer – not for bringing him into existence, but for putting him in a tiny cage, and for slaughtering him. But if Dickey's life were worth living, he does not have a complaint against the diner. The diner could not have saved him – they could only have prevented him from existing, and that would not be in Dickey's interest.

The implication of this argument is that *it is not wrong for harm-based reasons to cause someone to exist who is then abused by someone else, provided that their lives are worth living, and there was no alternative act which would have caused them to exist with a better life.* This is the typical position of the diner in relation to the animals their purchase affects.

This principle, I claim, is plausible even when applied to uncontroversial full-moral-status human beings – a test failed spectacularly by almost every existing defence of omnivorism, including Lomasky's appeal to existence-dependence.

Imagine that you are at a convenience store considering whether to purchase a piece of chocolate, and a perfectly reliable Oracle who is waiting behind you in line taps you on the shoulder and gives you the following information: if you purchase the chocolate, through a complicated series of events triggered by your purchase, a boy named Theodore will be born as a third child to a family which will abuse him in many morally terrible ways, and he will die young. Nevertheless, Theodore will have a life worth living – he will be glad to be alive, though he will regret the way he is treated. If you do not purchase the chocolate, Theodore will never be born. For good measure, the Oracle shows you a sketch of what Theodore will look like, a hint of sadness in his eyes.

Intuitively, it is not wrong to buy the chocolate. And this intuition is not grounded in mere causal distance or in the fact that the chain leading to Theodore's suffering is mediated through the agency of another. If the Oracle had told you that Theodore will exist no matter what, and purchasing the chocolate would, through an equally complicated series of events, cause him to suffer additional abuse that he would not otherwise face, it would be clearly wrong to buy the chocolate.

The diner's defence amounts to the claim that things are not any worse if the chocolate were

pork, and if the Oracle had shown you a picture of Dickey instead.

5. The Non-Identity Problem

The diner's defence shows, I believe, that as far as harm is concerned, it is not wrong to consume meat. But a little more must be said before we can be sure *which* meat we are permitted to eat, both at the level of theory and in practice.

We have already discussed why the argument would be limited to meat produced in conditions where life is worth living. Are there any additional constraints? Features often attributed to humane farming practices, for example in what McMahan [2008] calls 'benign carnivorism' – the animals having lives better than they would in the wild, living a large portion of their lifespan, being killed painlessly, and being replaced, are not necessary. Even if the Oracle told you Theodore's life did not meet conditions of this sort, it would not seem wrong to buy the chocolate. This suggests that perhaps the lives-worth-living bar is the *only* condition. Ultimately, I think this is basically correct. But we should not be too quick on this point.

So far, I have been assuming that the diner's choice is simple: order the meat or order vegetarian. If there is only one kind of animal and it is raised with a life worth living, then I think we have shown that it is permissible to order it. But in the real world we have many options when it comes to ordering meat. There are different kinds of animal, and even within a species, some animals are raised in better conditions than others. Are we obligated to eat animals that are raised in the *best* conditions?

This question turns out to be rather complicated. Suppose I can order two kinds of steak, one of which was raised in conditions which only barely pass the worth-living bar, and the other in conditions that are quite a bit better than that. My purchase, we'll assume, would raise demand for whichever kind of beef I order, and bring into existence (or risk bringing into existence) more cows in those conditions. But at the level of individual identity those cows would not overlap – the cows I bring into existence by ordering the more humane beef are entirely different individuals than the cows I bring

into existence by ordering the less humane beef. On the face of it then, by person-affecting criteria, it seems that no matter which beef I order, none of the cows will have a complaint against my act.

This case is closely related to a set of puzzles Parfit [1984] famously discusses under the label of the *Non-Identity Problem*. Parfit argues that if a woman has the choice to conceive one of two children, both of which would have lives worth living but one of which would be born with a severe defect that negatively impacted its quality of life, it would be wrong for her to conceive the one with the defect. Similar reasoning would suggest that I must order the more humane beef. But Parfit's response is controversial. In his thorough book-length treatment of the topic, David Boonin [2008] argues that accepting Parfit's judgment about the conception case ultimately leads to even more counterintuitive implications elsewhere, and so that even in that case, it is permissible to bring about the existence of the child with a worse life.¹⁶

I find Boonin's argument compelling, but the issue is too complicated to fully explore here. I will say only a few brief things in favour of the view that ordering either dish is permissible. The case of the two steaks is different from the conception case in ways that may be morally relevant. First, ordering the two steaks is not (in the normal case) exclusive, in the way choosing to conceive the children is. Ordering one of the steaks doesn't prevent you from ordering the other. This makes the steaks case more like two independent decisions – first, whether to bring one cow into existence, and second, whether to bring the other. Even Parfit does not think it is wrong to bring the child with the worse life into existence, if your choice is between that and no child, and that it is not wrong to bring no child into existence, even if your alternative is the child with a very good future. If this sequence of decisions on the part of a mother (far enough apart temporally that the choices do not exclude each other) does not seem wrong, then it suggests ordering either steak is permissible.

Second, in the case of the mother's choice, her intention is to bring someone into the world, while in the case of the diner's choice, the creation of a new being is merely a foreseen potential side-effect. It is not implausible to think that bringing the best possible life into being might matter more

¹⁶ It is worth noting that it is consistent with the claim that it is *permissible* to bring into existence the child with the worse life that it would be *better* to bring into existence the other child.

when bringing someone into existence is the intention.

Finally, in cases that more closely resemble the choice between steaks, I think intuition does not clearly condemn making a choice which leads to someone worse-off coming into existence. Suppose that at the convenience store, there is a whole array of snacks in front of you, delicious in different ways and to different degrees. The Oracle has conveniently attached to each a dossier of a child that will someday be born if you purchase that particular snack. Some of them have miserable lives, some of them have lives which are fantastic, and some of them have lives which are merely mediocre. It is wrong, it seems to me, to buy any of the snacks attached to a dossier of a child with a life that is worse than nothing. But it does not seem clearly wrong to order any of the snacks where the child has a decent life, even if you leave snacks on the shelf with better dossiers attached.

6. Which Meat Can We Eat?

If the reader does not share my judgment about the two steaks, then it would seem that the omnivore must take care to eat only the most humanely produced meat available to them. Otherwise, it should be permissible to eat any meat as long as it is produced under conditions that are better than nonexistence. What does this mean in practical terms?

Unfortunately, settling the question of which animals are raised in conditions better than nonexistence is too difficult for us to attempt here, partly because judging the value of lives is a difficult philosophical and empirical challenge of its own, and partly because farming practices differ between countries and across time. In the case of humans, we might get at least a first pass at such judgments by simply asking a person how they feel about their own life, or whether they would rather never have been born. But we cannot get these answers from nonhuman animals, and what I have to say here must be speculative. At least one thing should be relatively uncontroversial – at the time of this writing, many common factory farming practices do not produce animals with lives worth living.

Egg-laying hens in the United States, for example, frequently live their lives in battery cages with as little as 67 square inches of space per hen. They do not have space to spread their wings or to satisfy

powerful natural behavioural urges such as dust bathing and spend most of their lives in a state of extreme stress. A significant portion of these die during confinement from asphyxiation or dehydration. Because these conditions frequently drive them to extreme and injurious behaviour such as pecking at themselves or other hens, much of their beak is often severed without anaesthesia. This kind of life does not seem remotely worth living.

In other intensive farming practices, it is less clear whether animals have lives worth living. Perhaps the best case can be made for beef. The typical beef cow in the United States spends approximately half of its life (around seven months) grazing, and most of the rest in a feedlot being fattened for slaughter. They suffer significantly during the processes of branding, dehorning, and castration, and during travel and slaughter itself. Because of the significant time spent out of confinement and away from acute sources of stress, a case might be made that such lives pass the bar. Emphasizing recent positive advances in the treatment of cattle, this is the judgment of animal scientist and humane farming proponent Temple Grandin [Johnson 2015].¹⁷

To be safe, however, a conscientious omnivore convinced by the diner's defence should look into practices in their own country, and limit their consumption to meat, preferably beef, that is produced in conditions better than those at the most intensive farms.

7. Conclusion

The diner's defence is not meant to be a conclusive argument against moral vegetarianism. It is a defence against the most common type of reasoning against eating meat – namely, that it contributes to the harms and unnecessary slaughter inflicted upon farmed animals. Almost all moral vegetarian philosophers accept some argument of this sort. But there are other arguments against eating meat that the diner's defence cannot undermine, since they do not appeal to harms inflicted as a result of purchasing meat. One possible lesson to draw, then, is that vegetarians should turn their attention to these other arguments.

¹⁷ For discussion of the conditions of different animals under modern factory farming practices, see e.g. Rollin [1995].

It may be that eating meat is wrong because of its effect on the *environment* [Henning 2011]. It may be wrong because of its *symbolic* value – that being a vegetarian is a way of ‘standing for the good.’ [Cuneo 2015] It may be wrong because it amounts to a kind of *complicity* in the wrongdoing of the farmers [Driver 2015]. It may be wrong because it is *benefiting from* wrongdoing [Curnutt 1997]. It may be wrong because it is *rewarding* wrongdoing. It may be wrong because eating meat while knowing that it was produced immorally is a failure of *virtue* [Nobis 2002].

It is notable that the diner's defence leaves many of these open precisely because it does not defend the practice of eating meat by defending the practice of farming it. These arguments appeal to moral reasons that are more controversial than our moral reasons not to cause harm, which every plausible view accepts. And the reasons involved are, I think, generally weaker than harm-based reasons. For example, perhaps enjoying videos of cruel pranks is wrong because it fails to stand for the good, or benefits from wrongdoing, or manifests some vice. But it is not *as* wrong as committing a cruel prank. So even if ultimately, one of these other arguments is successful, the diner's defence might show that the actions of omnivores are *less* morally bad than we otherwise might think, and that we should be sceptical of the more extreme claims of moral vegetarians – that eating meat is as bad as torturing puppies, for example [Norcross 2004], or setting kittens on fire [McPherson 2014].

I want to end with a note of caution. I have presented the argument in this paper because I think it is correct. But I am also worried that it is dangerous. There is a risk that ameliorating our sense of guilt for eating meat can make us complacent about the great wrongs that are inflicted upon animals in factory farms. This would be a failure to appreciate the limited scope of the diner's defence. I once thought that eating meat was indefensible, and now I am not so sure. But I remain convinced as much as ever that the treatment of animals in the last century will be remembered as one of the darkest blots in our shared moral history.¹⁸

¹⁸ For invaluable feedback on earlier drafts of this paper, I would like to especially thank Bob Beddor, Ben Blumson, Alexander Dietz, Preston Greene, Joe Horton, Qu Hsueh Ming, Michael Pelczar, and Neil Sinhababu, as well as two anonymous referees for this journal.

References

- Boonin, David 2014. *The Non-Identity Problem and the Ethics of Future People*, Oxford: Oxford University Press.
- Bramble, Ben 2015. The Case Against Meat, in *The Moral Complexities of Eating Meat*, ed. Ben Bramble and Bob Fisher, Oxford: Oxford University Press: 135-52.
- Budolfson, Mark 2018. The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response, *Philosophical Studies*, advance online publication, URL = <<https://doi.org/10.1007/s11098-018-1087-6>>
- Cuneo, Terence 2015. Conscientious Omnivorism, in *Philosophy Comes to Dinner*, ed. Andrew Chignell, Terence Cuneo, and Matthew C. Halteman, New York: Routledge: 21-38.
- Curnutt, Jordan 1997. A New Argument for Vegetarianism, *Journal of Social Philosophy* 28/3: 153-72.
- Driver, Julia 2015. Individual Consumption and Moral Complicity, in *The Moral Complexities of Eating Meat*, ed. Ben Bramble and Bob Fisher, Oxford: Oxford University Press: 67-79.
- Harman, Elizabeth 2011. The Moral Significance of Animal Pain and Animal Death, in *The Oxford Handbook of Animal Ethics*, ed. Tom L. Beauchamp and R.G. Frey, Oxford: Oxford University Press: 726-737.
- Harris, J.R. and Galvin, R. 2012. 'Pass the Cocoamone, Please': Causal Impotence, Opportunistic Vegetarianism and Act-Utilitarianism, *Ethics, Policy, and Environment* 15/3: 368-83.
- Henning, Brian G. 2011. Standing in Livestock's 'Long Shadow': The Ethics of Eating Meat on a Small Planet, *Ethics and the Environment* 16/2: 63-93.
- Johnson, Nathanael 2015. Interview with Temple Grandin. URL = <<https://grist.org/food/temple-grandin-digs-in-on-the-practical-side-of-what-animals-want/>>
- Kagan, Shelly 2011. Do I Make a Difference?, *Philosophy and Public Affairs* 39/2: 105-41.
- Lomasky, Loren 2013. Is it Wrong to Eat Animals?, *Social Philosophy and Policy* 30/1-2: 177-200.
- McMahan, Jeff 2008. Eating Animals the Nice Way, *Daedalus* 137/1: 66-76.
- McPherson, Tristram 2014. A Case for Ethical Veganism, *Journal of Moral Philosophy* 11/6: 677-703.
- Narveson, Jan 1987. On a Case for Animal Rights, *The Monist* 70/1: 31-49.

- Nobis, Nathan 2002. Vegetarianism and Virtue: Does Consequentialism Demand Too Little? *Social Theory and Practice* 28/1: 135-56.
- Norcross, Alistair 2004. Puppies, Pigs, and People: Eating Meat and Marginal Cases, *Philosophical Perspectives* 18/1: 229–45.
- Norwood, F. Bailey and Lusk, Jayson L. 2011. *Compassion by the Pound: The Economics of Farm Animal Welfare*, Oxford: Oxford University Press.
- Parfit, Derek 1984. *Reasons and Persons*, Oxford: Oxford University Press.
- Parfit, Derek 2017. Future People, the Non-Identity Problem, and Person-Affecting Principles, *Philosophy and Public Affairs* 45/2: 118-57.
- Rollin, Bernard E. 1995. *Farm Animal Welfare: Social, Bioethical, and Research Issues*, Malden, MA: Wiley-Blackwell.
- Singer, Peter 1980. Utilitarianism and Vegetarianism, *Philosophy and Public Affairs* 9/4: 325-37.
- Singer, Peter 1993. *Practical Ethics*, Cambridge: Cambridge University Press.
- Stephen, Leslie 1896. *Social Rights and Duties: Addresses to Ethical Societies*, London: S. Sonnenschein & Company, ltd.
- Visak, Tatjana 2013. *Killing Happy Animals: Explorations in Utilitarian Ethics*, London: Palgrave Macmillan UK.