Journal of Ethics and Social Philosophy
VOL. 24, NO. 2 · MARCH 2023
https://doi.org/10.26556/jesp.v24i2.2334

ETHICAL VEGANISM AND FREE RIDING

Jacob Barrett and Sarah Zoe Raskoff

THE CREATION of animal products on factory farms causes animals a tremendous amount of pain and suffering. Learning about the severity and extent of this suffering often leads people to change their dietary choices or at least to feel some moral pressure to do so. Many therefore seem implicitly to accept that when deciding what to purchase and eat, animal suffering makes a significant moral difference. Moral philosophers often make this thought explicit, attempting to ground an obligation to go vegan in the horrifying consequences of factory farming.

It turns out, however, to be surprisingly difficult to explain how animal suffering generates any reasons to alter our dietary choices, much less an obligation to go vegan. The standard argument is that we should go vegan to reduce animal suffering.[1] But this argument faces a challenge: thanks to the size and structure of the animal agriculture industry, any individual's consumption decisions are overwhelmingly unlikely to make a difference. *Producing* animal products may be harmful and wrong, but the effects of any individual's *consumption* decisions are insignificant. Going vegan, in other words, is causally inefficacious.[2] The reduction of animal suffering cannot ground an obligation to go vegan.

The causal inefficacy objection poses a serious challenge to ethical veganism: the view that we have a moral obligation to refrain from purchasing and consuming animal products. Ethical vegans have carefully outlined the conditions of animals on factory farms, the suffering they experience, and the moral significance of reducing this suffering. But the causal inefficacy objection threatens to cut off the rationale for ethical veganism from animal suffering altogether, rendering irrelevant ethical vegans' persuasive arguments on these points. And, in this paper, we argue that common replies to the causal inefficacy objection

1 E.g., Singer, *Animal Liberation* and "Utilitarianism and Vegetarianism."

2 Budolfson, "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response" and "Is It Wrong to Eat Meat from Factory Farms?"; Fischer, *The Ethics of Eating Animals*, ch. 4; Nefsky, "Consumer Choice and Collective Impact" and "Fairness, Participation, and the Real Problem of Collective Harm"; Shahar, *Why It's OK to Eat Meat*, ch. 4.

are unsatisfactory. Attempts to show that individual vegans are indeed causally efficacious are unsuccessful or at best inconclusive. And arguments that appeal to factors like complicity not only face substantive difficulties but also fail to accommodate the moral significance of *reducing* animal suffering. A stronger argument for ethical veganism would acknowledge that the case for going vegan would be weaker if vegans *as a group* made no difference to reducing animal suffering either. It would tie our obligation to go vegan to the fact that vegans collectively reduce animal suffering, even if no individual vegan makes a difference.

Specifically, we believe that the best response to the causal inefficacy objection relies on the wrongness of free riding. The basic idea is this. As a group, individuals who abstain from animal products create a morally important good: a large reduction in animal suffering. If one recognizes this much yet consumes animal products, one is free riding on vegans. One is recognizing the value of their goal, recognizing that the group makes a significant difference to achieving it, and yet making an exception of oneself by free riding on, rather than participating in, its production. This is wrong because free riding is wrong. And it remains wrong even for those who do not recognize the value of this goal because morality does not let one off the hook so easily: one cannot escape an obligation to go vegan simply by not caring about animal suffering. We are not only obligated to produce morally important goods through our own actions but are also obligated to participate in, rather than free ride on, their collective production. The latter obligation explains why we should go vegan.

We begin by sketching the standard argument for ethical veganism. We then explain the causal inefficacy objection and why we find existing rejoinders inadequate. From here, we develop our anti–free riding argument and consider several objections that lead us to qualify but not abandon our conclusion. The upshot is that even if one settles several controversial issues in ways that make trouble for our argument, there at the very least remain strong reasons for most people to purchase significantly fewer inhumanely raised animal products.

To be clear, our goal is only to examine the connection between animal suffering and our reasons to go vegan or otherwise change our behavior. Parallel considerations apply to other negative consequences of animal agriculture. But nothing we say here bears on the plausibility of grounding reasons to go vegan in something other than the consequences of animal agriculture—for example, in the idea that eating animals is disrespectful.

1. THE STANDARD ARGUMENT

There is overwhelming evidence that factory farming causes animals immense suffering. Animals are kept in horrifying conditions, have their bodies mutilated

without anesthetic, are transported in overcrowded trucks in which they crush one another, and are slaughtered, often painfully and in a state of extreme fear. All this is well documented elsewhere; we spare readers the gruesome details.[3] But it is worth noting the scale. In the United States alone, about 9 billion chickens, 120 million pigs, and 30 million cows are slaughtered each year. Of these, 98.2 percent of chickens raised for eggs, 99.9 percent of chickens raised for meat, 98.3 percent of pigs, and 70.4 percent of cows are raised on factory farms.[4] The severity and scale of suffering on factory farms is staggering.

Pointing to this suffering is the first step in the standard argument for ethical veganism. The second step is to claim that this suffering massively outweighs any compensating benefits of factory farming, such as the pleasure of eating meat or associated cultural experiences.[5] This step, too, is familiar, but in brief, those who deny it face two challenges. First, it is notoriously difficult to find grounds for believing that animals' interests matter less than humans' interests that do not implausibly commit us to thinking that some humans' interests matter less than others'.[6] Second, even if humans' interests matter more, it is implausible that they matter *so much* more that the value we derive from animal products outweighs the extent of animal suffering. Perhaps if humans suffered as much as animals on factory farms do, this suffering would be even more abhorrent. But that is not the relevant comparison. The extent of animal suffering remains abhorrent, and extremely morally bad, even after factoring in any countervailing benefits.[7]

We believe that the first two steps of the standard argument are successful and assume as much here. Our focus is on the final step, which attempts to derive an obligation to go vegan or otherwise change our dietary behavior from the moral horrors of factory farming. To keep things simple, we focus on veganism for now and take up relevant differences later. According to this final

3 See, e.g., Singer, *Animal Liberation*; Foer, *Eating Animals*.

4 Jacy Reese Anthis, "US Factory Farming Estimates," Sentience Institute, April 11, 2019, https://www.sentienceinstitute.org/us-factory-farming-estimates. The statistic for cows refers to those raised in concentrated animal feeding operations as defined by the US Environmental Protection Agency. However, an anonymous referee points out that beef cattle in such operations lead much of their lives in better conditions than our description suggests. For now, we treat all animal products equivalently, but we return to relevant differences later.

5 On such benefits, see Lomasky, "Is It Wrong to Eat Animals?" See also Cohen, "The Case for the Use of Animals in Biomedical Research."

6 Singer, *Animal Liberation* and "Utilitarianism and Vegetarianism"; though see Kagan, *How to Count Animals*; Setiya, "Humanism."

7 DeGrazia, "Moral Vegetarianism from a Very Broad Basis"; Gill, "On Eating Animals"; Kagan, *How to Count Animals*, 5.

step, the connection between veganism and animal suffering is straightforward. Factory farms operate to meet consumer demand for animal products. Going vegan is an effective and relatively low-cost way to reduce demand and so to reduce animal suffering. And since we are obligated to reduce suffering when we can do so at relatively low cost, we are therefore obligated to go vegan.

2. THE CAUSAL INEFFICACY OBJECTION

Unfortunately, the standard argument is too simple. The trouble is that the animal agriculture industry's size and structure appear to render it insensitive to individual consumption decisions. For example, recall that nine billion chickens are slaughtered a year in the United States, working out to over twenty-four million chickens a day. This suggests that if you are at a restaurant deciding between chicken or tofu, your choice will not harm any chickens. The chickens in the restaurant are already dead. And going forward, it is not like ordering tofu sends a signal directly to a factory farmer who says, "Oh! I better produce one less chicken next month!" Instead, the restaurant purchases chickens in bulk, from a distributor who purchases in bulk, from a processor who purchases in bulk, and so on, all the way back to a factory farm. At no point in this supply chain are decisions fine grained enough to reflect individual choices—again, in the United States, *twenty-four million* chickens are slaughtered a day. Rather, reductions in demand only trigger reductions in supply when a series of thresholds is met: enough consumers must refrain from purchasing chickens from enough restaurants and stores that enough restaurants and stores reduce their orders from enough distributors, and so on, that enough processors reduce their orders from factory farms that those farms produce fewer animals. And one choice to order tofu is, unfortunately, not going to trigger all these thresholds. This is the causal inefficacy objection to ethical veganism.

This objection is powerful. The standard argument says that we should go vegan because doing so will reduce animal suffering and the cost is relatively small. But the causal inefficacy objection suggests that going vegan has no such benefit and so is not worth even a small cost. This appears to let omnivores off the hook: they can maintain that factory farming is awful, curse their causal inefficacy, and eat animal products with a clean conscience. Animals should not suffer, but abstaining from animal products does not reduce this suffering. So why should someone have to suffer through tofu when they so much prefer chicken?

Some ethical vegans are unconvinced. They argue that even though no individual decision is likely to make a difference, there must be some number of dietary choices that is large enough to do so—that is, to trigger the aforementioned series of thresholds. For example, even though no choice to abstain

from a whole chicken dinner is likely to save a chicken, perhaps every hundred thousand fewer chickens sold will result in one hundred thousand fewer chickens being produced. In that case, since you have no idea how far you are from triggering the relevant thresholds, you should assign a probability of 1/100,000 to your abstention from chicken triggering the thresholds and saving one hundred thousand chickens. And, according to expected value theory, saving one hundred thousand chickens with a 1/100,000 probability is exactly as good as saving one chicken with certainty. Thus, the causal inefficacy objection appears defused. Refraining from chicken is extremely unlikely to save any animals—but when it does, it saves a huge number of them. And this is enough to render one's choice efficacious.[8]

This expected value argument may seem to seal the deal for ethical veganism. But, again, things are not so simple. There are two basic worries—one calling into question the collective impact of large groups of vegans, another the impact of individual vegans.[9] The first worry is that even if one hundred thousand decisions to refrain from chicken would reduce the number of chickens on factory farms, they are unlikely to reduce the number by one hundred thousand. Instead, they will cause the price of chicken to drop, which will cause some who would not have otherwise bought chicken to do so—at least insofar as farmers are willing to produce chickens at the lower price.

Some proponents of the causal inefficacy objection suggest, on the basis of anecdotal evidence, that such "inelasticities" massively decrease the collective efficacy of vegans.[10] However, the food economists Norwood and Lusk estimate that this effect is slight: in the United States, for a sufficiently large number n, if n choices are made not to buy a chicken, $0.76n$ fewer chickens will be produced.[11] If this is roughly correct, then inelasticities cannot plausibly undermine the obligation to go vegan, since the expected benefit of saving roughly 0.76 chickens from suffering is still very significant relative to the associated cost. So the standard argument for ethical veganism withstands the worry that large groups of vegans are causally inefficacious. At least over the long run,

8 Norcross, "Puppies, Pigs, and People," 232–33; Singer, "Utilitarianism and Vegetarianism," 335–36; Kagan, "Do I Make a Difference?"

9 Compare Nefsky, "Consumer Choice and Collective Impact," 273. A third worry is that expected value theory treats small probabilities of large value inappropriately: a 1/100,000 probability of saving one hundred thousand chickens is not as good as saving one chicken with certainty. We set this aside here.

10 Budolfson, "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response," 1718, and "Is It Wrong to Eat Meat from Factory Farms?" 86–89; Nefsky, "Consumer Choice and Collective Impact," 273.

11 Norwood and Lusk, *Compassion, by the Pound*, 223. Their values for other animal products range from 0.56 to 0.91.

the animal agriculture industry is sensitive to macrolevel market trends—for example, to large numbers of individuals refraining from chicken.[12]

This brings us to the individual impact worry. Even if n decisions to refrain from chicken would trigger a series of thresholds that saves $0.76n$ chickens, it does not follow that a single decision to abstain has a $1/n$ probability of doing so: the *expected impact* of one decision need not equal the *average impact* of n decisions. The animal agriculture industry involves a long supply chain from the farm to the table. At each link in this chain there is considerable *slack*—a margin for error or tolerance for waste. Thanks to slack, one decision to abstain from chicken may have much less than $1/n$ the expected impact of n decisions, and so much less expected value than saving 0.76 chickens.[13] Consider a stylized case.

Suppose a grocery store decides how many chickens to order each month based on its sales the previous month. It orders chickens in quantities of 500, and typically sells about 9,750 chickens a month. So, for some time, it has been ordering 10,000 chickens a month. The store is willing to tolerate some waste and so will only reduce its order from 10,000 to 9,500 if it sells fewer than 9,500 the previous month. In this case, reducing chickens purchased from the store by 500 is guaranteed to result in the store purchasing 500 fewer chickens the following month. But what is the probability a single decision to refrain from chicken makes this difference? Is it 1/500, as proponents of the expected value argument assume?

Well, that depends. The probability your decision makes a difference would be 1/500 if your choice had an equal probability of reducing monthly demand to 9,499 as it did of reducing demand to any other level. But this need not be so. For example, suppose you know trends will hold up: the store will always sell between 9,600 and 9,900 chickens a month. Then, the probability you reduce next month's order by 500 is not 1/500, but zero. Or suppose you know the store will sell between 9,600 and 9,900 chickens in a month unless a shock occurs, in which case any level is equally probable, and there is a 99 percent probability no shock occurs. Then, you have a 1/500 probability of causing the store to order 500 fewer chickens *if* a shock occurs, but only a 1/50,000 probability of doing so overall.

These examples demonstrate that expected impact can come apart from average impact and so defeat any *a priori* argument that the two must converge:

12 McMullen and Halteman, "Against Inefficacy Objections"; compare Hedden, "Consequentialism and Collective Action," 536.

13 Budolfson, "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response" and "Is It Wrong to Eat Meat from Factory Farms?" See also Nefsky, "Consumer Choice and Collective Impact"; Fischer, *The Ethics of Eating Animals*, ch. 4.

it is possible that vegans are on average impactful, yet the expected impact of one vegan is very low. The question therefore becomes whether, empirically, we find this divergence. Skeptics argue that we do not because consumers lack the crucial information about thresholds and trends generating the results in the above examples. Absent such information, slack makes no difference to expected impact. Suppose you are in the same store, but you have no idea where thresholds lie or what consumer trends are like. Then, you should estimate that you are just as likely to reduce demand to a threshold value as you are to reduce it to any other level. So, if thresholds occur every n choices, the probability you hit a threshold is $1/n$. And this remains true even if you know there is a lot of slack in the animal agriculture industry. More slack implies that the distance between thresholds is larger, not that your expected impact is lower: it implies that n is larger, not that you have less than a $1/n$ probability of hitting a threshold.[14]

But this is too quick. Even absent information about consumer trends or the location of thresholds , consumers might nevertheless know that the probability they reduce demand to precisely some threshold level is lower than $1/n$ because there is some correlation between thresholds and consumer demand.[15] For example, suppose the grocery store decided how many chickens to purchase this month based on its projection of how many it would sell, and it will only change its future orders if this projection proves far off. Then, if the grocery store's projection is fairly reliable—larger errors in its projection are much less probable than smaller ones—the probability this projection is far off will be much lower than the probability it is approximately right. So you are *not* just as likely to reduce demand to the far-off threshold level that makes the store change its future orders as you are to reduce it to any other level. The probability you hit the threshold is much less than $1/n$.

Proponents of the causal inefficacy objection argue that this is relevantly analogous to the position real consumers find themselves in.[16] Grocery stores lose customers who see empty cases, so they have an incentive to overpurchase animal products.[17] This is partly why they produce so much waste—in the United States, for example, a lower-end estimate suggests that grocery stores

14 Hedden, "Consequentialism and Collective Action," 537–39; compare McMullen and Halteman, "Against Inefficacy Objections," 99–100.

15 Hedden notes this possibility but dismisses it as empirically unlikely ("Consequentialism and Collective Action," 539n17).

16 See Fischer, *The Ethics of Eating Animals*, ch. 4; Shahar, *Why It's OK to Eat Meat*, ch. 4. They draw on Budolfson, "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response" and "Is it Wrong to Eat Meat from Factory Farms?"

17 Fischer, *The Ethics of Eating Animals*, 59; Shahar, *Why It's OK to Eat Meat*, 99.

throw away 4.5 percent of fresh meat.[18] Since demand fluctuates from month to month, grocery stores have a further incentive to tolerate significant variation in sales without adjusting their orders, rather than trying to fine-tune their orders and risk underpurchasing, leaving their customers unhappy.[19] So they have strong incentives to adopt ordering strategies much like in our example. They make projections with some margin of error, such that only large divergences from their projections (signaling a change in market trends), but not minor fluctuations, lead them to change their plans. As our stylized case shows, this can be enough to reduce a consumer's expected impact.

Indeed, the problem may be worse in the real world because similar dynamics arise at each link of the supply chain—the effects of which can compound rapidly to drive down expected impact. Not only do grocery stores and restaurants tolerate waste and variations in demand when deciding how much to purchase from their distributor, so does their distributor when deciding how much to buy from their processor, and so on. Unsurprisingly, then, there is a huge amount of slack in real-world supply chains. For example, in North America and Oceania, a whopping 13 percent of meat initially produced on farms is wasted while working its way down the supply chain to the consumer.[20]

All this suggests that one's expected impact may be considerably lower than proponents of the expected value argument assume. Of course, just how low is hard to say: the failure of any *a priori* argument that expected impact equals average impact means that one's expected impact depends on controversial claims about messy empirical reality. Here, Norwood and Lusk's figures on elasticity are often cited as showing that individuals have a high expected impact.[21] But this is a misinterpretation, as these figures only refer to average impact, and Norwood's stated view is that "for all practical purposes, expected effects are impossible to determine."[22] Any confident assertion that a typical consumer's expected impact is high enough relative to the cost of going vegan to ground an obligation to do so therefore strikes us as overconfident. This either defeats

18 Buzby et al., "Supermarket Loss Estimates for Fresh Fruit, Vegetables, Meat, Poultry, and Seafood and Their Use in the ERS Loss-Adjusted Food Availability Data." See Fischer, *The Ethics of Eating Animals*, 59, for discussion.

19 Fischer, *The Ethics of Eating Animals*, 60.

20 Food and Agriculture Organization, *Global Food Losses and Food Waste*.

21 E.g., MacAskill, *Doing Good Better*, 88, 228.

22 Fischer, *The Ethics of Eating Animals*, 61 (personal correspondence with Norwood). The misinterpretation is understandable, since Norwood and Lusk do say that buying one unit less of chicken reduces production by 0.76 units (*Compassion, by the Pound*, 223). But (as Norwood confirms in the same personal correspondence with Fischer) their analysis only shows that buying n fewer units reduces production by $0.76n$ units *for large n*'s.

the standard argument for ethical veganism—which relies on the premise that individuals have a high expected impact—or, depending on where one places the burden of proof, at least leaves the defender and critic of ethical veganism at a stalemate, as far as the standard argument is concerned.[23] But we need not settle for a draw. Controversy about expected impact leaves untouched the (comparatively) uncontroversial claim that vegans have a huge collective impact on reducing animal suffering, both absolutely and relative to the cost of going vegan. Intuitively, this should be enough to ground an obligation to go vegan, regardless of whether *individual* vegans have a significant expected impact. Can some other account explain why?

3. TOWARD A SOLUTION

We began with the idea that the suffering of animals on factory farms somehow grounds an obligation to abstain from animal products. The simplest view is that we should go vegan because doing so reduces suffering, but this runs into the causal inefficacy objection. Many ethical vegans have therefore retreated to the idea that we should go vegan even if doing so fails to reduce animal suffering, because our dietary choices bear some other relation to animal suffering. A typical view appeals to complicity. The animal agriculture industry produces tremendous suffering. Participating in it makes one wrongfully complicit.

Standard objections to complicity views challenge their explanations of either why omnivores are complicit or why complicity is wrong. For example, suppose we understand complicity *causally*, in terms of an individual's expected impact on the maintenance of factory farming, or *expressively*, in terms of what eating animals conveys—say, approval of factory farming or callous disregard for animal suffering.[24] The causal inefficacy objection challenges the view that

23 Compare Fischer, *The Ethics of Eating Animals*, 62. For the same reason, attempts to revive the expected value argument by appeal to "indirect effects"—going vegan may cause others to go vegan (Norcross, "Puppies, Pigs, and People," 233; Almassi, "The Consequences of Individual Consumption," 404–7)—are unsatisfactory. Such effects are similarly difficult to determine, especially since they are not uniformly positive: vegans may turn others *off* veganism. Although we cannot find a quantitative estimate, our suspicion is that the typical consumer's expected indirect effects are (on balance) positive but small. The most plausible route to high indirect impact is indirect effects compounding, as when each vegan converts two others to veganism, who each converts two others, and so on. But this is not what we find: veganism, unfortunately, is not growing at an exponential rate (compare Nefsky, "Consumer Choice and Collective Impact," 271).

24 See, respectively, Lepora and Goodin, *On Complicity and Compromise*, ch. 3; and Driver, "Individual Consumption and Moral Complicity."

omnivores are causally complicit.[25] And those convinced by the causal inefficacy objection need not be expressively complicit either: they might abhor factory farming and feel for the suffering of animals, yet consume animal products because they believe doing so makes no difference.[26] Or suppose we understand complicity *extractively*, in terms of benefiting from wrongdoing.[27] Omnivores are plausibly extractively complicit, but there are many counterexamples to the view that benefiting from wrongdoing is wrong. To take just one: it does not seem wrong to benefit from medical procedures that were developed through wrongful experimentation.[28]

We are sympathetic to these objections, but rather than fortifying them against potential rejoinders, we focus on a more general problem. The puzzle the causal inefficacy objection raises is that while vegans collectively reduce animal suffering, it is hard to see why individuals are obligated to join in if no individual has a significant impact. An adequate solution to this puzzle should provide some way of connecting individuals' obligations to go vegan with the collective reduction of animal suffering, rather than rendering this reduction irrelevant. Complicity views cannot do this since they sever the connection between the obligation to go vegan and the reduction of animal suffering altogether: we would still be wrongfully complicit, on such views, even if no number of vegans could successfully reduce animal suffering. But vegans do collectively reduce animal suffering, and an argument for ethical veganism can and should take advantage of this fact.

The importance of tying our obligation to go vegan to vegans' collective impact comes out clearly when we consider another view about why eating some animal products is wrong even if individuals are causally inefficacious—the view that it is wrong to consume products whose production *essentially* involves harm, regardless of our causal relation to that production.[29] On this view, it is permissible to eat animal products—say, milk—that *can* be produced painlessly, even if their actual production involves suffering. But it is wrong to eat products that are essentially harmful in the sense that they cannot be produced without suffering, even if refraining from them does not reduce this harm.

This account is problematic. Suppose a new technology allows us to produce meat using a device that extracts flesh from animals before immediately

25 Budolfson, "Is It Wrong to Eat Meat from Factory Farms?" 92, and "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response," 1713.

26 Fischer, *The Ethics of Eating Animals*, 96.

27 McPherson, "Why I Am a Vegan (and You Should Be One Too)."

28 Nefsky, "Consumer Choice and Collective Impact," 277.

29 Budolfson, "Is It Wrong to Eat Meat from Factory Farms?" 94–97.

repairing their bodies without causing any suffering. Since the technology renders harm no longer essential to the production of meat, its development should, on this account, result in the consumption of meat transforming from impermissible to permissible—even if it is never used. This is implausible. Actually using this futuristic technology to produce meat would render eating that meat morally unproblematic (at least on grounds of its connection to animal suffering), but the mere existence of this technology lacks such significance. The lesson is that our moral obligations depend on the actual harm and suffering caused by animal production, rather than this harm's essentiality. More carefully, they depend on the extent to which vegans collectively *reduce* this harm.

To drive this point home, suppose that producing bacon causes the same amount of suffering as producing chicken wings, but (due to differences in their supply chains) no number of individuals refraining from chicken wings would make any difference to this suffering, while relatively small numbers refraining from bacon would make a significant difference. Then, intuitively, even if no individual makes a difference, we have stronger reasons to refrain from bacon than to refrain from wings. And, crucially, this holds regardless of which *essentially* involves more harm or makes us more complicit.

A final view that accommodates this insight holds that individuals have reasons to help bring about valuable outcomes, such as the reduction of animal suffering, even when individual participation makes no difference.[30] This view relies on a nonstandard definition of helping, on which an individual helps bring about an outcome when their action plays a *nonsuperfluous* causal role in its production. An action counts as nonsuperfluous when it is possible both for it to be part of the cause of that outcome and for the outcome to fail to come about because not enough people perform actions of that type. So we have reasons to help in such cases, and our reasons to help more impactful collectives are stronger.[31]

Unlike the other views on offer, this view successfully connects reasons to go vegan with the reduction of animal suffering: individuals have reasons to *help* reduce animal suffering. But we worry that it ultimately leaves such reasons unexplained. It is uncontroversial that we have reasons to help bring about outcomes if "helping" is interpreted in its everyday sense, which involves making a causal difference. But the view in question uses "helping" (and associated notions like "nonsuperfluous") as terms of art. And reasons to help in this special sense do not automatically inherit the credentials of reasons to help in the everyday sense. Absent some further explanation of why we have

30 Nefsky, "How You Can Help, without Making a Difference."

31 Nefsky, "How You Can Help, without Making a Difference," 2764.

such reasons, then, this view appears to relocate rather than solve our puzzle: Why do individuals have reasons to "help" if no instance of helping makes a difference? A more satisfying and complete view would provide this further explanation.[32] We develop such a view now.

4. THE ANTI–FREE RIDING ARGUMENT

We have seen that an adequate reply to the causal inefficacy objection should hold on to the idea that our obligation to go vegan is somehow related to the reduction of animal suffering. Specifically, it should tie our obligation to go vegan to the collective impact of vegans, even if no individual makes a difference. Thankfully, morality provides us with just the connection we need. In many contexts, it is wrong to free ride on the collective production of important goods. Since the reduction of animal suffering is one such good that vegans collectively produce, those who consume animal products free ride, and this is what makes their consumption decisions wrong. We have moral obligations not only to reduce animal suffering through our own actions but also to participate in, rather than free ride on, collective endeavors that have this impact. The causal inefficacy objection suggests that going vegan may not be an effective way to discharge the former obligation, but it leaves the latter untouched. This explains why we should go vegan.

There is much disagreement about how exactly to formulate the moral requirement not to free ride, but there is also widespread agreement that some such requirement applies in paradigm cases. It is a crucial feature of these cases that free riding is wrong even if individual instances of free riding make no difference. For example, many explain why we should pay to ride public transportation—rather than literally riding for free—by appeal to the wrongness of free riding. Although our own measly fare is unlikely to affect the public transit system, the system would cease to function if nobody paid to ride. Since anybody could equally help themselves to the justification that *their* participation makes no difference, allowing that this provides a genuine exemption from paying would preclude us from the good of public transportation. Morality helps us overcome this problem by denying that this is a genuine exemption and by demanding that we pay our fare even if doing so makes no difference.

32 A related problem of incompleteness arises for McPherson, who suggests that there may be reasons to participate in collectively beneficial social patterns and to abstain from collectively detrimental ones ("The Puzzle of Inefficacy," 240–42). Such reasons might solve our puzzle, but we need a further explanation of why they exist. (To be fair to McPherson, he does not purport to defend this view, but merely gestures to it as having the right general shape to solve our puzzle.)

This, if nothing else, is what seems right about Kantian approaches to morality: we ought not to make an exception of ourselves, at least in cases where anyone else could make the same exception, and where granting this exception would preclude the production of an important good.[33]

It is true that paradigmatic cases of wrongful free riding involve free riding on the collective production of what economists call a *public* good: a (nonexcludable, nonrivalrous) good that benefits the free rider rather than some third party. It is because *I* benefit from riding public transit that I should contribute my fare; those who do not ride public transit are not guilty of free riding. This might seem to render the case of reducing animal suffering disanalogous from standard cases of wrongful free riding: it is animals rather than vegans who benefit from the collective action of vegans. We address this worry in two steps.

The first step is to note that those who acknowledge the great importance of reducing animal suffering do derive value from the actions of vegans. To put things, again, as economists might: those who value the reduction of animal suffering would be willing to pay some monetary cost to reduce it, and so benefit from its reduction, at least in the broad sense of "benefit" relevant here. Indeed, there are many familiar examples of public goods whose production you should not free ride on even though they benefit you, not in the narrow sense that they further your self-interest, but in the broad sense that they further what you care about or prefer to achieve. Common examples include poverty reduction and herd immunity against illnesses that do not threaten you (assuming you care about others' welfare and health). In fact, certain governments, such as the United Kingdom's, explicitly identify animal welfare as a public good given widespread preferences for improved animal welfare.[34] It is similarly common to treat animal welfare as a public good in economic analyses.[35]

For those who recognize the great value of reducing animal suffering on factory farms, then, the anti–free riding argument succeeds. Reducing animal suffering is an important public good, and it is a familiar feature of such goods that we should not free ride on their collective production even if our own contribution makes no difference.

This conclusion is already significant. Dialectically, the causal inefficacy objection is mounted by those who acknowledge the great importance of reducing animal suffering and so would acknowledge an obligation to go vegan if individuals were causally efficacious. Our argument implies that anyone in

33 This preliminary gloss on the wrongness of free riding owes much to Cullity, "Moral Free Riding" and "Public Goods and Fairness."

34 Coe and Finlay, *The Agriculture Act 2020*.

35 E.g., Norwood and Lusk, *Compassion, by the Pound*, ch. 10.

this camp should go vegan, so it responds to the causal inefficacy objection in the dialectical context in which it typically appears.

But this may seem unsatisfying. The standard argument for ethical veganism defends the stronger conclusion that one has an obligation to go vegan regardless of whether one cares about animal suffering. And it seems odd to accept our weaker conclusion that those who care about animal suffering have this obligation, without also accepting the stronger one. After all, to claim otherwise is to allow that individuals can duck their moral obligations—in this case, not to purchase or consume animal products—simply by failing to care (or not caring enough) about what they morally ought to care about—in this case, animal suffering. Morality, however, should not let such callousness get us off the hook. If those who care about animal suffering have an obligation to go vegan, then so too should those who do not.

The second step of our argument bridges the gap between the weaker and the stronger conclusions. The most straightforward route appeals to the idea that callousness cannot extinguish obligations. The argument is simple. The first step of our anti–free riding argument establishes that those who care about animal suffering should go vegan. But, we now add, callous indifference to animal suffering cannot absolve one of such an obligation. So one must have an obligation to go vegan regardless of whether one cares about animal suffering. More generally, if callousness does not exempt one from moral obligations, then the existence of an obligation not to free ride on the collective production of goods one cares about implies an associated obligation not to free ride on the collective production of goods one morally ought to care about, in the sense that not caring about them would exhibit the moral failing of (perhaps among other things) callousness. In other words, it implies a general obligation not to free ride on the collective production of goods that we either care about or morally ought to care about. And this general obligation entails the stronger conclusion that one should go vegan regardless of whether one cares about animal suffering.

This anti-callousness argument is difficult to resist. Given the weaker conclusion that those who care about animal suffering have an obligation not to free ride on collective endeavors that reduce it, one can block it only by insisting that callousness can absolve one of this obligation. But this is implausible. In the individual case, it would be a nonstarter for someone to claim that they are not obligated to produce a morally important good simply because they do not care about it. And it seems equally bizarre to allow that simply not caring about a morally important good can absolve one of an obligation to participate in a collective endeavor to produce it. If callousness does not exempt us of obligations in individual cases, then it should not do so in collective cases either. "I do not care about the suffering of others" is no better a justification for failing

to participate in collective endeavors that reduce suffering than it is for failing to reduce suffering through one's own actions.

Our approach so far has been to argue from the widely acknowledged obligation not to free ride on the collective production of public goods to a similar obligation not to free ride on the collective production of morally important goods. We now strengthen our case for the existence of this latter obligation—first, by drawing out a counterintuitive implication of rejecting it and, second, by noting that most moral theories converge on it.

First, in the absence of this obligation, morality would include a counterintuitive loophole absolving individuals of their obligations to participate in collective action in cases where no individual makes a difference, even though (i) the group is successfully producing a morally important good, (ii) the group is only able to produce this good because its members participate despite their individual inefficacy, and (iii) individuals would be obligated to participate if the group produced the same benefit but through a different causal mechanism that rendered each individual causally efficacious. But we submit that morality does not have loopholes. Much as in the case of riding public transit, morality does not grant us permission to free ride when that permission would be equally available to everyone, and where granting it would preclude the creation of an important good.

To make the counterintuitiveness of this loophole vivid, note that if it existed, morally motivated agents could be manipulated by bad actors in a peculiar way. Suppose that advances in technology reduce slack to the point that every dietary choice makes a significant causal difference. In this world, morally motivated individuals (who otherwise prefer omnivorism) go vegan to reduce animal suffering by their individual actions, driving factory farms to the brink of viability. Realizing their error, factory farmers hatch a devious plot: they will exploit the loophole by reintroducing slack into the system, rendering individual dietary choices again inefficacious. This has no effect on the collective impact agents have on reducing animal suffering, but assuming the loophole in question exists, it removes any obligation to go vegan. The factory farmers' plot succeeds. Morally motivated agents resume consuming animal products and so cease reducing animal suffering.[36]

We think it clear that something has gone wrong in this story. Morality should not include a loophole allowing bad actors to manipulate morally motivated agents in this way: changes to individual efficacy should not lead morally motivated agents to stop participating in morally important collective action

36 We adapt this thought experiment from Budolfson, "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response," 1713n1.

that is efficacious at the group level. An obligation not to free ride is exactly the right shape to plug this loophole. So morality must include one.[37]

Second, although our argument has proceeded at the intuitive level, we may strengthen it by noting that most moral theories converge in endorsing an obligation not to free ride on the collective production of morally important goods. Most obviously, pluralistic deontologists may accept this obligation as a foundational duty, or as one grounded in a requirement of fairness. Indeed, the most influential and well-developed account of free riding holds that free riding on the collective production of both public and morally important goods is unfair, and hence wrong, for the same reason: roughly, because both involve exempting oneself from an obligation in a context where the production of a good requires individuals not to exempt themselves in this way.[38]

Other moral theories accommodate the same obligation.[39] For example, rule-consequentialists will ground it in its good consequences, because individuals adhering to it produces better outcomes. Contractualists will agree, holding that no one can reasonably reject this obligation since its rejection would place larger burdens on those who would otherwise benefit from the production of such goods.[40] Virtue theorists will ground an obligation not to free ride in the virtue, say, of being *cooperative*, where cooperative people are willing to join in morally important collective endeavors.[41] It is only orthodox act-consequentialists who cannot accommodate this obligation, but that is unsurprising.

An obligation not to free ride on the collective production of morally important goods generally, and so on the collective reduction of animal suffering specifically, therefore stands on firm ground. Shortly, we will consider some challenges to our claim that this obligation implies that we must go vegan,

37 Our claim here is not that bad actors can *never* manipulate morally motivated agents but merely that the possibility of the particular sort of manipulation in the above example—involving a reduction of individual efficacy but no change to collective efficacy or to anything else—is especially implausible. Thanks to an anonymous referee for raising this issue.

38 Cullity, "Public Goods and Fairness" and *The Moral Demands of Affluence*, 62–65. More carefully, Cullity holds that I unfairly free ride when I rely on others to do their parts in a collective imperative to bring about some good without doing my own part. The grounds of the collective imperative differ depending on whether the good is public or morally important, but free riding on each is unfair in the same way. (Cullity reserves the term "free riding" for cases of public goods, but this difference is merely terminological.)

39 Compare Brennan, "Polluting the Polls," 540.

40 On some versions of contractualism, burdens on animals do not count, rendering such an obligation inapplicable to the case at hand. We take such theories to be implausible, at least insofar as they are meant to provide a complete theory of morality.

41 Cullity, *Concern, Respect, and Cooperation*, ch. 3.

which will require us to qualify our view. But first we clarify our position by addressing three unsuccessful objections.

5. CLARIFYING THE POSITION

The first objection comes from those sympathetic to our conclusion yet skeptical that our argument provides the right explanation of why it is wrong to purchase and consume animal products. The anti–free riding argument suggests that we are obligated to go vegan because it is wrong to free ride on *vegans*. Animals, here, seem oddly missing from the story.

It is true that our explanation of why one should go vegan is that, otherwise, one is free riding on vegans. But the anti–free riding argument still accommodates the intuitive sense in which a concern for animals ultimately grounds this obligation, because there remains the further question of why it is wrong to free ride on vegans. And our answer is that this is wrong because vegans are collectively producing a morally important good: the reduction of animal suffering. If animal suffering were morally unimportant, or if vegans were failing collectively to reduce it, the anti–free riding argument would not succeed. Our explanation of why it is wrong to purchase and consume animal products therefore makes essential reference to the moral importance of reducing animal suffering. Although it would be theoretically neater to claim that individuals should go vegan to reduce animal suffering by their own actions, the causal inefficacy objection calls this into doubt. Our account therefore provides exactly what we set out to provide: a plausible way of grounding an obligation to go vegan in the moral importance of reducing animal suffering. Animals are not missing from the story; they take center stage.

A second objection concerns cost. We only have obligations not to free ride, the thought goes, when we can do so at relatively low cost. We are happy to grant this. But this is unsurprising; proponents of ethical veganism nearly universally acknowledge that the obligation to go vegan does not apply to those for whom veganism would be especially costly. With the exception of those for whom veganism poses a serious economic burden or those who have a relevant medical condition, however, this limitation does little work. Indeed, despite common misconceptions, a vegan diet is cheaper for most people in high- and upper-middle-income countries.[42] It is healthier than typical omnivorous diets.[43] And vegans do not lead worse lives in general, at least as far as

42 E.g., Springmann et al., "The Global and Regional Costs of Healthy and Sustainable Dietary Patterns."

43 See Garrett, "Utilitarianism, Vegetarianism, and Human Health," for an argument that the expected value of going vegetarian is therefore significantly positive (with a survey of the

their (self-reported) subjective well-being is concerned.[44] Real hardship may justify omnivorism, just as it may justify those in severe poverty sneaking onto public transit and riding for free. But most people reading this paper are lucky enough not to have this justification available.

We return to cost in the next section, but for now consider a third objection: that the strength of our reason not to free ride depends on how many others are participating. The worry is that if almost no one is participating in the production of some good, then one's reason to participate is weaker and may not generate an obligation. And since vegans are a small minority, this might seem to undermine the anti–free riding argument.

There is something to this thought. In many cases, the fact that very few are participating in collective action extinguishes our obligation not to free ride. Consider an analogous case. Early in the COVID-19 pandemic, social distancing was encouraged as a method for slowing the spread of the virus. Certain models, however, suggested that social distancing is only beneficial when nearly universally practiced.[45] If this is correct, then under conditions where social distancing was nowhere near, and was never likely to reach, this level, individuals did not seem to have an obligation to social distance (grounded in the wrong of free riding). Since individuals engaged in social distancing were failing to reduce virus transmission, others were not wrongfully free riding by failing to do so. The lesson is that one cannot free ride on the production of a good if too few are participating to produce that good, or perhaps if one does not anticipate that enough will eventually participate.

The same example can, however, be repurposed to more constructive ends. Other models of social distancing suggest that even though no individual makes a significant difference, the more who engage in social distancing the better: 50 percent of people social distancing reduces virus transmission more than 40 percent, which reduces it more than 30 percent, and so on. On these models, social distancing produces not a binary (or single-step) good but a multistep good. The good is not all or nothing, only kicking in at, say, 70 percent of people social distancing. Rather, small numbers of people social distancing produce the good of reduced virus transmission to some degree, and increases in the number of people social distancing result in this good being produced to a greater degree if those increases are sufficiently large.

---

empirical evidence at notes 24 and 25).

44 See Pfeiler and Elgoff, "Do Vegetarians Feel Bad?" for an attempt to measure this while controlling for various factors, and Iguacel et al., "Vegetarianism and Veganism Compared with Mental Health and Cognitive Outcomes," for a meta-analysis.

45 Chang et al., "Modelling Transmission and Control of the COVID-19 Pandemic in Australia."

If these latter models are correct, then, intuitively, we did have an obligation to social distance grounded in the wrong of free riding, even when, say, only 25 percent of people were doing so (at least assuming the collective benefits were worth the cost). What matters is not the total number or proportion of people who are participating in the production of a good, but whether enough are participating to collectively make a significant difference—and perhaps whether further increases in participation will, at some point, pass some further threshold or step that results in even greater production of the good.

As we have seen, the reduction of animal suffering is a multistep good in the same way. Even if no individual vegan makes a difference, large enough increases in decisions to abstain from animal products do. So the fact that there are few vegans does not imply that we lack an obligation not to free ride on them, so long as they are genuinely reducing animal suffering. Our obligation not to free ride on the collective production of multistep goods is not extinguished merely because a small minority are participating. It is only extinguished when too few are participating to make any difference at all.[46]

According to a 2018 Gallup poll, approximately 3 percent of individuals in the United States (where veganism is not especially popular) self-identify as vegans.[47] This strongly suggests that vegans are indeed making a collective difference: as we have argued, although the supply chain for animal products may be insensitive to individual decisions, it is highly sensitive to macrolevel market trends. As further evidence of this, consider how many more vegan products are available at restaurants and grocery stores than there used to be. It follows by the anti–free riding argument that it is wrong to free ride on vegans' reduction of animal suffering, even if 97 percent of people in the United States are guilty of doing so.

46 Contra Nefsky ("Fairness, Participation, and the Real Problem of Collective Harm," 255), we see it as a benefit, not a cost, that the anti–free riding argument implies no obligation to go vegan when not enough are doing it to make a difference. But we stress that this does not further imply that we are never obligated to initiate novel forms of collective action. We may very well have such an obligation for some other reason—say, due to the expected impact (via the unusually large "indirect effects") of starting a new movement. But, crucially, the actions required to initiate collective action often differ from those required not to free ride on ongoing collective endeavors—for example, going vegan may, but need not, be an effective way to start a movement to reduce animal suffering (compare Cullity, *Concern, Respect, and Cooperation*, ch. 3). So it is an advantage of our account that it does not treat the two cases identically. (And, if one disagrees, one may modify our account to say that we also have an obligation not to free ride on *potential* morally important goods.)

47 R. J. Reinhart, "Snapshot: Few Americans Vegetarian or Vegan," Gallup, August 1, 2018, https://news.gallup.com/poll/238328/snapshot-few-americans-vegetarian-vegan.aspx.

6. WIGGLING OUT OF VEGANISM

This brings us to an important issue we have bracketed so far. We have framed our discussion as an argument that we have an obligation to go vegan, rather than to modify our choices in related ways. But have we really earned this conclusion? There are four worries to consider.

The first worry is that our argument relies on the moral significance of animal suffering and therefore cannot explain why purchasing and consuming animal products from humane, suffering-free farms is objectionable. As we have noted, however, the overwhelming majority of animal products in the United States and much of the world comes from factory farms. And there are reasons to doubt whether allegedly humane farms are genuinely humane: a 2012 report from the Animal Welfare Institute, for example, suggested that production methods on US chicken farms certified by the US Department of Agriculture as "superior" for animal welfare "are not materially different from conventional production methods" and that standards are very laxly enforced.[48] Many ethical vegans therefore argue that, in practice, there is hardly any difference between buying whatever animal products one pleases and only buying animal products that are produced in allegedly humane environments. But there may be exceptions, and some certifications do seem more reliable than others. Indeed, the Animal Welfare Institute has more recently published a useful guide of which animal welfare food labels are trustworthy and which are misleading.[49]

Here, there remain further questions about the morality of raising and slaughtering animals, even painlessly, and so about the morality of buying animal products from genuinely humane farms. On views on which painlessly killing animals is bad (say, because doing so violates rights), we can construct a structurally parallel anti–free riding argument against purchasing and consuming animal products from humane farms. This argument, however, will be weaker in force since those who purchase meat from humane farms are free riding only on the reduction of animal deaths, rather than on the reduction of both deaths and suffering. But whether it is bad for animals to be painlessly slaughtered is much less obvious than whether it is bad for animals to suffer, and we will not investigate the question further here.

A second worry is that our argument runs together purchasing and consuming animal products, but it might seem only to establish an obligation not to purchase them. It might leave dumpster diving or eating dumpster-bound

48 Mathews, "Humanewashed," 1.

49 Animal Welfare Institute, "A Consumer's Guide to Food Labels and Animal Welfare."

leftovers permissible, since even widespread engagement in such activities does not increase demand for animal products.

We are unsure about this. Perhaps individuals who publicly abstain from eating free animal products significantly contribute, as a group, to the reduction of animal suffering through their influence on others' dietary choices or by helping to dismantle a broader ideology of "carnism" that upholds factory farming.[50] Or perhaps individuals who abstain from consuming animal products in private are less likely to purchase animal products in the future and so collectively have a larger impact.[51] Again, we are unsure. These are the sorts of empirical questions the anti–free riding argument requires us to attend to. Like the standard argument, it makes our reasons to refrain from animal products contingent on facts about the empirical consequences of our dietary choices—only it focuses us on the effects of groups rather than individuals and requires us not to free ride on efficacious groups.

This brings us to the third way our argument may fall short. Vegans are not the only ones who collectively reduce animal suffering. So too, say, do vegetarians. At first glance, our argument suggests no principled case for going vegetarian rather than vegan since dairy cows and chickens raised for eggs suffer on factory farms. Nevertheless, we may consider vegetarianism one strategy of reducing one's use of animal products, akin to other strategies like adhering to Meatless Mondays, Veganuary, or just trying to eat less meat. And it is not obvious that our argument establishes an obligation to go vegan rather than merely to reduce.

Now, the anti–free riding argument does suggest that we have stronger reasons to go vegan than merely to reduce: reducetarians, after all, free ride on vegans, whereas the reverse is not true.[52] So in the absence of weighty counter-

50 See Joy, *Why We Love Dogs, Eat Pigs, and Wear Cows*, ch. 2. Compare John and Sebo ("Consequentialism and Nonhuman Animals," 575–76), who argue against eating humanely raised meat on similar grounds.

51 Compare Almassi, "The Consequences of Individual Consumption," 407–8; John and Sebo, "Consequentialism and Nonhuman Animals," 572–74.

52 We here assume that reducetarians form a larger collective of which vegans are a subset. This raises thorny questions about collective membership conditions, since, thanks to the causal inefficacy objection, we cannot simply count someone as a member when they make a difference to the outcome a collective produces (Nefsky, "Fairness, Participation, and the Real Problem of Collective Harm"). In general, the key to solving this issue is to identify members of a collective by their contribution to the same "underlying dimension" in which each *does* make a difference (say, the number of animal products purchased), where large enough changes in this underlying dimension cause changes to the morally important dimension (say, the amount of animal suffering) (Wieland and Oeveren, "Participation and Superfluity"). However, we do not defend any particular account of the relevant dimension here and instead rely on intuitive judgments about collective membership.

vailing reasons to reduce rather than to go vegan (deriving from the greater cost of veganism), the anti–free riding argument suggests that we must go vegan. And it might seem like such reasons are unavailable, since, as we have argued, veganism is not excessively costly, at least for most healthy and affluent people.

Yet this is too quick. The claim we defended above is that the difference in cost between going vegan and failing to participate in the collective reduction of animal suffering is insufficient to outweigh our reasons to participate, which derive their strength from the large collective impact of vegans on reducing animal suffering. But this leaves open the possibility that veganism is excessively costly when compared to some forms of reducing, even for healthy and affluent people, since veganism might be significantly more costly than reducing yet have little additional collective impact.

For example, suppose that someone finds it somewhat costly to go vegan but almost costless to eat largely vegan while allowing themselves some wiggle room, say, in certain social contexts or when traveling. If the difference in the collective impact of strict veganism and wiggly veganism is low but the difference in cost between the two is relatively significant, then perhaps the person is justified in being wiggly vegan. Similarly, if the difference in collective impact between veganism and vegetarianism is low and one finds it much less costly to eliminate meat but not eggs and cheese from one's diet, then perhaps merely going vegetarian is justified.

We find it plausible that some are justified in going less than strictly vegan for such reasons. But this concession is not as significant as it might seem. If one can eliminate the bulk of the cost of going vegan by building in some wiggle room, then doing so may be justified, but this provides no license for failing to go vegan outside narrowly tailored cases. For example, suppose one finds veganism costly because one's social life is organized around meals with nonvegan family and friends. Then, even granting that the significance of such social costs permits one wiggle room in social situations, this provides no justification for eating animal products when alone or in social contexts where vegan options are available.

So while it is hard to say precisely how wiggly any individual's veganism may be, we have not yet found a challenge to our claim that most of us are obligated to significantly reduce our consumption of animal products. Our discussion has assumed, however, that strict vegans are collectively more impactful than reducetarians, such that we have stronger reasons to be strict. But the opposite may hold if reducetarians more often inspire others to join the cause, say, because strict veganism tends to turn others off.[53] If so, reducetarianism may be morally preferable.

53 Thanks to an anonymous referee for this point.

Fourth and finally, consider the possibility that one might discharge one's obligation not to free ride without altering one's dietary choices at all. Suppose someone generously donates their money or time to animal welfare charities or activism. Does the anti–free riding argument imply that they must *also* go vegan (or reducetarian)? Or have they already done everything required of them to discharge their obligation not to free ride?

In fact, the worry is more general. We have defended an obligation to alter one's dietary choices as an implication of a general obligation not to free ride on the collective production of morally important goods. But an obligation not to free ride on *all* relevant collective endeavors might seem overly demanding, even if participating in any particular endeavor is not. This suggests that we might have latitude about not only how to participate in the collective reduction of animal suffering but also whether to participate in this or some other collective endeavor when discharging our general obligation not to free ride.[54] Put another way, it might seem implausible that each group can obligate *everyone* to devote themselves to its cause, simply by initiating effective collective action in its pursuit.[55]

This raises deep questions about morality's demandingness that we cannot fully address here. But there are, broadly speaking, three relevant sorts of views. The first is an extremely demanding view on which moral demands *iterate*. If particular demands are not too costly, then we must meet each of them, even if the demands add up to something highly costly in aggregate. On this view, we are obligated to go vegan (or reducetarian) since doing so, considered on its own, is not too costly. So too are we obligated to participate in every other way and in every other cause, so long as each additional form of participation, considered on its own, is not too costly.

But less demanding conceptions of morality can also ground an obligation to go vegan or reducetarian. Consider a moderately demanding conception that factors in the *aggregate* cost of various demands, but with further provisos. For example, perhaps morality cannot demand so much of us, in aggregate, that we are left unable to lead a flourishing and autonomous life, with various personal projects and commitments. But it demands that when choosing between projects, we opt for morally preferable ones when the resulting package of projects does not leave us significantly worse off or conflict with core commitments (including to our loved ones). This is rough, and the devil is in the details.[56] But

54 Shahar, *Why It's OK to Eat Meat*, ch. 5. Compare Fischer, *The Ethics of Eating Animals*, ch. 9.

55 Thanks to an anonymous referee for articulating the point this way.

56 For one attempt to spell out such a moderately demanding view, see Cullity (*The Moral Demands of Affluence*, ch. 9), from whom we have also adapted the *iterative* versus *aggregative* distinction.

views of this type will generally imply that we must participate in the collective reduction of animal suffering. Since one must eat regardless of how else one spends one's time, and since we have seen that veganism is cheaper and so leaves one with more resources than omnivorism without undermining one's (subjective) well-being, adding veganism (or reducetarianism) to one's portfolio of projects is clearly compatible with leading a good life full of other projects and commitments.[57] This represents a disanalogy with many other forms of activism, the addition of which would threaten to swallow up one's time and resources and so one's other projects and commitments.

Finally, consider mildly demanding conceptions of morality on which one must pay some aggregate cost, and that is all. On such views, we are not obligated to go vegan or reducetarian so long as we meet the relevant cost threshold in other domains of life. However, even here it is implausible that we have *absolute* latitude or equally strong reasons to participate in any way and in any cause. For example, we intuitively have weaker reasons to donate to less important causes and in less (collectively) effective ways. And our reasons to participate in the collective reduction of animal suffering are plausibly very strong, even if not uniquely so. In the first place, veganism—or at least reducetarianism—is unusually collectively cost effective, in the sense that it is low cost at the individual level but produces a massive benefit at the collective level. There are few other cases, if any, where individuals can sacrifice so little to collectively do so much.[58] Second, when it comes to reducing animal suffering, there is, so to speak, no neutral option. It is not as if the only alternative to joining in a collective endeavor, say, to reduce domestic violence is to (individually or collectively) produce domestic violence. Most of us do neither. But in the case of reducing animal suffering, neutrality is not an option. Some collectively produce animal suffering; others collectively reduce it. One must pick a side.

57 Shahar raises similar considerations but dismisses them on the grounds that they are only relevant on an extremely demanding conception of morality (*Why It's OK to Eat Meat*, 109–13). In so doing, he appears to overlook the possibility of a moderately demanding view.

58 Shahar is skeptical that veganism produces large collective benefits and denies that it is low cost because it causes social friction, involves "treating every meal as a weighty ethical decision," deprives one of gustatory pleasure, and involves transition costs, like relearning how to cook (*Why It's OK to Eat Meat*, 116–19). However, we take our earlier discussion of elasticity and average impact to support high collective impact: the average decision to abstain from chicken, recall, saves 0.76 chickens, at least according to Norwood and Lusk (*Compassion, by the Pound*, 223). (Shahar approvingly cites this point [*Why It's OK to Eat Meat*, 107n14].) Further, as we have noted, one can, if necessary, reduce social friction by allowing certain exceptions, and as we will discuss shortly, one can reduce the cost of moral deliberation by adopting simple rules. This leaves transition costs and pleasure. But such costs, though they may loom large in the minds of omnivores, are small in comparison to the long-term (collective) benefits of maintaining vegan or reducetarian diets.

Both considerations suggest that our reasons to participate in the reduction of animal suffering are uncommonly strong.

So, on extremely and moderately demanding conceptions of morality, we are obligated to go vegan or reducetarian; on mildly demanding conceptions, we have strong reasons to do so. Or, at the very least, we have strong reasons to avoid purchasing factory-farmed products. Whether we should also avoid *consuming* animal products or purchasing *suffering-free* animal products—and how strict about this we should be—depends on how we answer the difficult empirical and normative questions we have raised in this section but cannot hope to resolve here. Instead, we end with three pragmatic considerations.

The first is that we generally have stronger reasons to avoid animal products for which similar reductions of demand would trigger greater reductions in suffering, and this can yield unintuitive results. For example, avoiding eggs plausibly has a larger collective impact than avoiding beef or pork since chickens raised for eggs are more abused and less efficient than beef cattle or pigs (as measured in animals per calorie produced), and since the supply chain for chicken and eggs is more elastic.[59] The second is that it is typically easier to consistently follow simple, rigid rules than vague or highly complex ones.[60] This does not bear on how strictly one should abstain from animal products, but it does suggest that a policy of, say, never purchasing animal products at grocery stores, going vegetarian, or eliminating chicken and eggs may be better than a policy of "eating less meat."[61] Adopting rigid rules that do not require deliberation on a case-by-case basis may also eliminate certain costs of reducing consumption, since many find it costly to treat what they previously saw as the "morally free" zone of dietary choices as a domain that now requires moral deliberation.[62] The third is that given our tendency to manifest a self-serving bias when engaging in moral reasoning, it may be that we should err on the side of adopting stricter policies. Those who enjoy eating animal products can easily convince themselves that there is nothing wrong with doing so, and

59 See MacAskill ("Effective Reducetarianism," 70), who argues that the three "most effective way[s] to reduce animal suffering [are] to stop eating chicken, then eggs, then pork." The issue is complicated by other negative consequences of animal agriculture, such as greenhouse gas emissions, which cattle disproportionately produce. But for an argument that animal agriculture has much larger welfare effects on animals than its emissions have on humans (drawing on climate economic models of the social cost of carbon), see Kuruc and McFadden, "Monetizing the Externalities of Animal Agriculture."

60 E.g., John and Sebo, "Consequentialism and Nonhuman Animals," 574–75.

61 Rothgerber finds that compared to vegetarians, "conscientious omnivores" report both greater difficulty following their diet and more frequent violations of it ("Can You Have Your Meat and Eat It Too?").

62 On this cost, see again Shahar, *Why It's OK to Eat Meat*, 116.

those of us trained in analytic philosophy are especially good at coming up with rationalizations of this choice.[63]

University of Oxford
jacob.barrett@philosophy.ox.ac.uk
sarah.raskoff@philosophy.ox.ac.uk

REFERENCES

Almassi, Ben. "The Consequences of Individual Consumption: A Defence of Threshold Arguments for Vegetarianism and Consumer Ethics." *Journal of Applied Philosophy* 28, no. 4 (November 2011): 396–411.

Animal Welfare Institute. "A Consumer's Guide to Food Labels and Animal Welfare." Washington, DC: Animal Welfare Institute, August 2021. https://awionline.org/sites/default/files/uploads/documents/21FoodLabelGuideFull.pdf.

Brennan, Jason. "Polluting the Polls: When Citizens Should Not Vote." *Australasian Journal of Philosophy* 87, no. 4 (December 2009): 535–49.

Budolfson, Mark Bryant. "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response." *Philosophical Studies* 176, no. 7 (July 2019): 1711–24.

———. "Is It Wrong to Eat Meat from Factory Farms? If So, Why?" In *The Moral Complexities of Eating Meat*, edited by Ben Bramble and Bob Fischer, 80–98. Oxford: Oxford University Press, 2015.

Buzby, Jean C., Hodan Farah Wells, Bruce Axtman, and Jana Mickey. "Supermarket Loss Estimates for Fresh Fruit, Vegetables, Meat, Poultry, and Seafood and Their Use in the ERS Loss-Adjusted Food Availability Data." *EIB*-44. US Department of Agriculture, Economic Research Service. March 2009.

63 The authors would like to thank Rhys Borchert, Mark Budolfson, David DeGrazia, Jamie Draper, Christopher Howard, Mario Juarez-Garcia, Joseph Millum, Andreas T. Schmidt, Danny Shahar, Hannah Tierney, Alice van't Hoff, David Wendler, and two anonymous referees for helpful comments and discussion. Thanks also to members of the Philosophy Work in Progress Group at the University of Arizona, the Davis Area Group in Ethics and Related Subjects at the University of California, Davis, and the Centre for the Study of Social Justice Seminar Series at the University of Oxford. This research was supported in part by the National Institutes of Health Clinical Center and the Intramural Research Program. The views expressed in this article are the authors' own. They do not represent the position or policy of the National Institutes of Health, the Department of Health and Human Services, or the United States government.

Chang, Sheryl L., Nathan Harding, Cameron Zacreson, Oliver M. Cliff, and Mikhail Prokopenko. "Modelling Transmission and Control of the COVID-19 Pandemic in Australia." *Nature Communications* 11 (November 2020): 19393–96.

Coe, Sarah, and Jonathan Finlay. *The Agriculture Act 2020*. Briefing Paper no. CBP 8702. London: House of Commons Library, 2020. http://researchbriefings.files.parliament.uk/documents/CBP-8702/CBP-8702.pdf.

Cohen, Carl. "The Case for the Use of Animals in Biomedical Research." *New England Journal of Medicine* 315, no. 14 (October 1986): 865–69.

Cullity, Garrett. *Concern, Respect, and Cooperation.* Oxford: Oxford University Press, 2018.

———. *The Moral Demands of Affluence*. Oxford: Oxford University Press, 2004.

———. "Moral Free Riding." *Philosophy and Public Affairs* 24, no. 1 (January 1995): 3–34.

———. "Public Goods and Fairness." *Australasian Journal of Philosophy* 86, no. 1 (March 2008): 1–21.

DeGrazia, David. "Moral Vegetarianism from a Very Broad Basis." *Journal of Moral Philosophy* 6, no. 2 (March 2009): 143–65.

Driver, Julia. "Individual Consumption and Moral Complicity." In *The Moral Complexities of Eating Meat,* edited by Ben Bramble and Bob Fischer, 67–79. Oxford: Oxford University Press, 2015.

Fischer, Bob. *The Ethics of Eating Animals: Usually Bad, Sometimes Wrong, Often Permissible.* New York: Routledge, 2020.

Foer, Jonathan Safran. *Eating Animals*. New York: Little, Brown and Company, 2009.

Food and Agriculture Organization. *Global Food Losses and Food Waste: Extent, Causes and Prevention*. Rome: Food and Agriculture Organization of the United Nations, 2011. http://www.fao.org/3/I2697E/i2697e.pdf.

Garrett, Jeremy R. "Utilitarianism, Vegetarianism, and Human Health: A Response to the Causal Impotence Objection." *Journal of Applied Philosophy* 24, no. 3 (August 2007): 223–37.

Gill, Michael B. "On Eating Animals." *Social Philosophy and Policy* 30, no. 1–2 (January 2013): 201–7.

Hedden, Brian. "Consequentialism and Collective Action." *Ethics* 130, no. 4 (July 2020): 530–54.

Iguacel, Isabel, Inge Huybrechts, Luis A. Moreno, and Nathalie Michels. "Vegetarianism and Veganism Compared with Mental Health and Cognitive Outcomes: A Systematic Review and Meta-analysis." *Nutrition Reviews* 79, no. 4 (April 2021): 361–81.

John, Tyler M., and Jeff Sebo. "Consequentialism and Nonhuman Animals." In *The Oxford Handbook of Consequentialism*, edited by Douglas W. Portmore, 564–91. Oxford: Oxford University Press, 2020.

Joy, Melanie. *Why We Love Dogs, Eat Pigs, and Wear Cows*. San Francisco: Conari Press, 2009.

Kagan, Shelly. "Do I Make a Difference?" *Philosophy and Public Affairs* 39, no. 2 (Spring 2011): 105–41.

———. *How to Count Animals, More or Less.* Oxford: Oxford University Press, 2019.

Kuruc, Kevin, and Jonathan McFadden. "Monetizing the Externalities of Animal Agriculture: Insights from an Inclusive Welfare Function." *Social Choice and Welfare* (forthcoming).

Lepora, Chiara, and Robert Goodin. *On Complicity and Compromise*. Oxford: Oxford University Press, 2013.

Lomasky, Loren. "Is It Wrong to Eat Animals?" *Social Philosophy and Policy* 30, no. 1–2 (January 2013): 177–200.

MacAskill, William. *Doing Good Better: How Effective Altruism Can Help You Help Others, Do Work That Matters, and Make Smarter Choices about Giving Back*. New York: Avery, 2015.

———. "Effective Reducetarianism." In *The Reducetarian Solution: How the Surprisingly Simple Act of Reducing the Amount of Meat in Your Diet Can Transform Your Health and the Planet*, edited by Brian Kateman, 69–71. New York: TarcherPerigee, 2017.

Mathews, Rachel. "Humanewashed: USDA Process Verified Program Misleads Consumers about Animal Welfare Marketing Claims." Washington, DC: Animal Welfare Institute, March 2012. https://awionline.org/sites/default/files/publication/digital_download/fa-humanewashedreportonusdapvp_0.pdf.

McMullen, Steven, and Matthew C. Halteman. "Against Inefficacy Objections: The Real Economic Impact of Individual Consumer Choices on Animal Agriculture." *Food Ethics* 2, no. 2–3 (May 2019): 93–110.

McPherson, Tristram. "The Puzzle of Inefficacy." In *Philosophy and Climate Change*, edited by Mark Budolfson, Tristram McPherson, and David Plunkett, 222–51. Oxford: Oxford University Press, 2021.

———. "Why I Am a Vegan (and You Should Be One Too)." In *Philosophy Comes to Dinner: Arguments about the Ethics of Eating*, edited by Andrew Chignell, Terence Cuneo, and Matthew C. Halteman, 73–91. New York: Routledge, 2016.

Nefsky, Julia. "Consumer Choice and Collective Impact." In *The Oxford Handbook of Food Ethics*, edited by Anne Barnhill, Mark Budolfson, and Tyler

Doggett, 267–86. Oxford: Oxford University Press, 2018.
———. "Fairness, Participation, and the Real Problem of Collective Harm." In *Oxford Studies in Normative Ethics*, vol. 5, edited by Mark Timmons, 245–71. Oxford: Oxford University Press, 2015.
———. "How You Can Help, without Making a Difference." *Philosophical Studies* 174, no. 11 (November 2017): 2743–67.
Norcross, Alastair. "Puppies, Pigs, and People: Eating Meat and Marginal Cases." *Philosophical Perspectives* 18, no. 1 (December 2004): 229–45.
Norwood, F. Bailey, and Jayson L. Lusk. *Compassion, by the Pound: The Economics of Farm Animal Welfare.* Oxford: Oxford University Press, 2011.
Pfeiler, Tamara M., and Boris Egloff. "Do Vegetarians Feel Bad? Examining the Association between Eating Vegetarian and Subjective Well-Being in Two Representative Samples." *Food Quality and Preference* 86 (December 2020): 104018.
Rothgerber, Hank. "Can You Have Your Meat and Eat It Too? Conscientious Omnivores, Vegetarians, and Adherence to Diet." *Appetite* 84 (January 2015): 196–203.
Setiya, Kieran. "Humanism." *Journal of the American Philosophical Association* 4, no. 4 (Winter 2018): 452–70.
Shahar, Dan C. *Why It's OK to Eat Meat*. New York: Routledge, 2022.
Singer, Peter. *Animal Liberation: A New Ethics for Our Treatment of Animals.* New York: Harper Collins, 1975.
———. "Utilitarianism and Vegetarianism." *Philosophy and Public Affairs* 9, no. 4 (Summer 2020): 325–37.
Springmann, Marco, Michael A. Clark, Mike Rayner, Peter Scarborough, and Patrick Webb. "The Global and Regional Costs of Healthy and Sustainable Dietary Patterns: A Modelling Study." *Lancet Planet Health* 5, no. 11 (November 2021): e797–e807.
Wieland, Jan Willem, and Rutger van Oeveren. "Participation and Superfluity." *Journal of Moral Philosophy* 17, no. 2 (April 2020): 163–87.