


ARTICLE

The Vegan's Dilemma

Donald W. Bruckner 

Penn State University, New Kensington
Corresponding author. E-mail: donald.bruckner@psu.edu

(First published online 6 April 2020)

Abstract

A common and convincing argument for the moral requirement of veganism is based on the widespread, severe, and unnecessary harm done to animals, the environment, and humans by the practices of animal agriculture. If this harm footprint argument succeeds in showing that producing and consuming animal products is morally impermissible, then parallel harm footprint arguments show that a vast array of modern practices are impermissible. On this first horn of the dilemma, by engaging in these practices, vegans are living immorally by their own lights. This first horn can be avoided by assuming that morality requires not minimizing harm, but only keeping the harm of our actions within some budget. On the second horn, however, we recognize that there are many ways of keeping our harm footprints within budget other than through our dietary choices. On the second horn of the vegan's dilemma, therefore, veganism is not a moral requirement.

1. Introduction

1.1. *The Harm Footprint Argument for Veganism*

There are powerful philosophical arguments for the conclusion that the production, consumption, and other use of the products of animal agriculture are wrong. One of the most common and convincing arguments draws this veganic conclusion on the basis of the widespread, severe, and unnecessary harm done to animals, the environment, and humans from the practices of producing and consuming animal products. To name but a few animal harms: farm animals suffer painful mutilation from castration, dehorning, and other procedures without anesthesia; they are forced to live in cramped quarters that prevent them from engaging in natural behaviors such as rooting for pigs or even spreading their wings for laying hens; they suffer rough handling, broken bones, and dehydration on the way to slaughter, which itself can be a horrific experience. On the environmental side, growing crops to feed animals whom we then eat is an incredibly inefficient way to produce nutrition in comparison to growing plants directly for human consumption. This results in harm to the environment for the vast waste of water, land, and fossil fuel. Fertilizers and manure harm the environment by polluting; ruminants pollute by producing a large volume of greenhouse gasses (GHGs). Finally, humans are harmed by animal agriculture due to harm to the environment as well the increased risk of many diseases from eating animal products.

© Cambridge University Press 2020

Many are also harmed in dangerous farm and slaughterhouse work. Most of this harm is unnecessary, for consuming just plant-based food is a viable and less harmful alternative. The usual argument claims that this massive, unnecessary harm footprint of producing and consuming products of animal agriculture implies that it is morally impermissible.

It is argued elsewhere that while many of the harms discussed above are standard by-products of conventional (“factory”) farming operations, most of the alleged harms are eliminated or dramatically reduced by small-scale animal agriculture that aims precisely at the elimination or reduction of these harms.¹ Further, plant agriculture contributes considerable harm to animals, the environment, and humans.² Still, while it is open to debate, a reasonable person could hold that even the harms of small-scale animal agriculture exceed those of plant agriculture, largely due to land use implications. So on the usual argument based on the harm footprint of animal agriculture, even the production and consumption of the products of small-scale animal agriculture are morally impermissible. For this reason, in what follows I just address animal agriculture in general, rather than providing distinct treatments of conventional and alternative forms.

The logic of the harm footprint argument is as follows. Suppose we have a practice P1 that produces massive harm to animals, the environment, and humans. Suppose that the harm that comes from P1 is unnecessary in the sense that another practice P2, which is a suitable alternative to P1, is available and produces less total harm. Then the harm footprint argument concludes that participation in or support of P1 is morally impermissible.

Notice that P2 itself may produce massive harm, yet still be permissible. The harm footprint argument does not claim that the production of massive harm is a sufficient condition for the impermissibility of a practice. There may be no available alternative that is less harmful. Plant agriculture is also massively harmful, but there is no suitable alternative. So it is permissible. If a suitable and less harmful practice were found, then the harm footprint argument would apply and plant agriculture would be impermissible. It is in this sense that the harm footprint argument requires harm minimization. This will be important in what follows.³

Instances of the harm minimization argument are now ubiquitous in popular and philosophical works. Consider two recent, explicit instances. David DeGrazia argues that “it is wrong to cause extensive, unnecessary harm” and that “[f]actory farming routinely causes extensive, unnecessary harm.”⁴ Since “vegetarian diets can be fully nutritious” alternatives, people like us who consume animal products but have access to such an alternative diet “have culinary lifestyles that are morally indefensible.”⁵ DeGrazia takes pains to point out that this requirement of moral vegetarianism comes from a very broad basis in the sense that his argument depends only on the very weak assumption that animals have some moral status.⁶

¹Donald W. Bruckner, Small-Scale Animal Agriculture, in *The Routledge Handbook of Animal Ethics*, ed. by Bob Fischer (New York: Routledge, 2019), pp. 198–210.

²See Bruckner, Small-Scale Animal Agriculture, pp. 200, 203, 205.

³While the assumption that we ought to minimize harm is usually implicit in harm-based arguments for veganism, sometimes it is explicit. For instance, Michael Allen Fox claims that “one of the guiding ideals of ecologically informed thinking is that we ought to minimize the harmful impact of our lives . . . on the biosphere” and that this principle of “minimizing harm . . . certainly seems to be about as basic a moral principle as can be imagined” (*Vegetarianism and Planetary Health, Ethics and the Environment*, 5 (2000), 163–74 (p. 166)).

⁴DeGrazia, Moral Vegetarianism from a Very Broad Basis, *Journal of Moral Philosophy*, 6 (2009), 143–65 (p. 155).

⁵DeGrazia, p. 150.

⁶DeGrazia, p. 147.

Hooley and Nobis lay out the harm footprint argument quite explicitly for any product the production of which “causes serious and unjustified harms,” where the product “can be avoided” while “not seriously harm[ing] the boycotting individual” because “there are readily available alternatives.” When these conditions are met for a product, “individuals are obligated not to purchase or use that product.”⁷ Just like DeGrazia, Hooley and Nobis stress that their argument appeals only to the very weak and uncontroversial assumption that “[i]t is wrong to cause serious harms . . . unless those harms are morally justified.”⁸

The fountainhead of this line of thinking about animal harm is, of course, Peter Singer’s *Animal Liberation*.⁹ Although very few instances of the harm minimization argument appeal to all three types of harm, the strongest version of the argument represented here does.¹⁰

A quick terminological note. ‘Veganism’ usually means abstaining from consumption and other uses of all animal products, including from hunting, fishing, and foraging. This article is concerned just with the possible moral requirement to abstain from products of animal agriculture. Given this scope and in order to avoid cumbersome language, hereafter I will just use ‘animal products’ as shorthand for ‘products of animal agriculture’ and ‘veganism’ to mean ‘abstinence from animal products’ in this narrower sense.

1.2. Causal efficacy, threshold chickens, and complicity

In order to situate the dialectic of this article, consider the *causal inefficacy objection* against harm footprint arguments for veganism. According to this objection, an individual’s dietary choices do not make a difference to the amount of harm that results from the production and consumption of animal products. After all, if Greta stops eating chicken, she will make absolutely no difference to the number of chickens produced. Markets simply are not sensitive to one individual’s consumption decisions. So while it would be wrong for Greta to cause massive, unnecessary harm to animals, humans, and the environment, her chicken-eating actions do not cause such harm. The objection concludes that since her dietary choices are causally inefficacious in this sense, it is not wrong for her to make dietary choices that support the practices of the chicken industry.

A first well-known response admits that most consumption decisions are causally inefficacious, but denies that this affects the strength of the harm footprint argument.¹¹ While

⁷Dan Hooley and Nathan Nobis, A Moral Argument for Vegetarianism, in *Philosophy Comes to Dinner: Arguments about the Ethics of Eating*, ed. by Andrew Chignell, Terence Cuneo, and Matthew C. Halteman (New York: Routledge, 2016), pp. 92–108 (p. 99).

⁸Hooley and Nobis, p. 93. For another argument in this spirit, which is supposed to appeal only to weak assumptions everyone accepts, see Mylan Engel, The Commonsense Case for Ethical Vegetarianism, *Between the Species*, 19 (2016), 2–31.

⁹Peter Singer, *Animal Liberation* (New York: Harper Collins, 1975).

¹⁰Other instances of this argument form include Tristram McPherson, Why I am a Vegan (And You Should Be One Too), in *Philosophy Comes to Dinner*, ed. by Chignell, Cuneo, and Halteman, pp. 73–91; Fox, Vegetarianism and Planetary Health; Ben Bramble, The Case Against Meat, in *The Moral Complexities of Eating Meat*, ed. by Ben Bramble and Bob Fischer (Oxford: Oxford University Press, 2016), pp. 135–50; Stuart Rachels, Vegetarianism, in *The Oxford Handbook of Animal Ethics*, ed. by Tom L. Beauchamp and R. G. Frey (Oxford: Oxford University Press, 2011), pp. 877–905; Peter Singer and James Mason, *The Way We Eat: Why Our Food Choices Matter* (Emmaus, PA: Rodale, 2006).

¹¹Versions of this response have been offered by Alastair Norcross, Puppies, Pigs, and People: Eating Meat and Marginal Cases, *Philosophical Perspectives*, 18 (2004), 229–45 (pp. 232–33); Peter Singer,

one consumer ceasing buying chicken will not cause a drop in chicken production, there is some number of consumers such that if that number of consumers ceases buying chicken, chicken production will drop in response by some large number of chickens. For illustration, suppose that the average consumer buys 10 chickens per year and that it would take 10,000 consumers ceasing buying chicken for chicken producers to respond by producing 100,000 fewer chickens per year. Then if Greta refrains from buying chicken this year, she has a 1 in 10,000 chance of saving 100,000 chickens. So the average consumer like Greta can expect that refraining from buying 10 chickens will save $\frac{100,000 \text{ chickens}}{10,000} = 10$ chickens. So even though it is highly unlikely that refraining from buying 10 chickens will save any chickens at all, refraining from buying 10 chickens is, as Norcross says, “morally and mathematically equivalent” to saving 10 chickens “from excruciating lives.”¹² Call this the *threshold chicken response* to harm footprint arguments. If the threshold chicken response succeeds, harm footprint arguments are successful even when an individual’s action is (almost certainly) causally inefficacious in the production or prevention of harm.

A second response to the inefficacy objection is similar to the threshold chicken response in admitting that the vast majority of our actions that participate in or support harmful or beneficial practices are causally inefficacious in the production or prevention of harm. This second response claims that, even so, participation in or support of that unnecessary and massively harmful practice is morally forbidden because it is wrong to be *complicit* in unnecessary and massively harmful practices.¹³ If Lucius Malfoy participates in the lynching of a muggle and then tries to excuse his behavior on the grounds that the unnecessary and massive harm would have been produced whether he participated or not, his behavior is no less wrong. Similarly, the *complicity response* to the inefficacy objection says, if one participates in or supports the massive and unnecessarily harmful practice of producing and consuming animal products, that participation or support is wrong on grounds of complicity even if one’s participation or refrain will make no difference to the production of harm.

There are other replies to the causal inefficacy objection and there are other objections against harm footprint arguments. I take the positions just outlined to be the leading contenders in the current debate. To keep things manageable, I limit discussion here to these positions.¹⁴ To review: When considering a harm footprint argument, there are three main positions the advocate of the harm footprint argument can take: (A) an individual’s actions are causally efficacious in the production or prevention of harm, so the harm footprint argument goes through; (B) an individual’s actions are not (usually) causally

Utilitarianism and Vegetarianism, *Philosophy and Public Affairs*, 9 (1980), 325–37 (pp. 335–36); Shelly Kagan, Do I Make a Difference?, *Philosophy and Public Affairs*, 39 (2011), 105–41 (p. 124); and DeGrazia, p. 158.

¹²Norcross, p. 233.

¹³Complicity arguments for veganism have been offered by Julia Driver, Individual Consumption and Moral Complicity, in *The Moral Complexities*, ed. by Bramble and Fischer, pp. 67–79; Adrienne M. Martin, Factory Farming and Consumer Complicity, in *Philosophy Comes to Dinner*, ed. by Chignell, Cuneo, and Halteman, pp. 203–14; and McPherson, Why I am a Vegan, pp. 83–85. For a sophisticated examination of complicity arguments, see Mark Budolfson, The Inefficacy Objection to Deontology: What it is, Why it is Important, and How to Respond to it, Version 2.6, unpublished manuscript, <<http://www.budolfson.com/papers/BudolfsonDeontologyFactoryFarms.pdf>> [accessed 19 December 2019].

¹⁴Another reply is skeptical and claims that one does not act wrongly by participating in a collectively harmful practice when one’s participation does not make a difference. Julia Nefsky surveys the state of the debate more broadly in her Collective Harm and the Inefficacy Problem, *Philosophy Compass*, 14 (2018), 1–17.

efficacious, but the threshold chicken response succeeds, so the harm footprint argument still goes through; and (C) an individual's actions are not causally efficacious, but the complicity response succeeds, so the harm footprint argument goes through again.

1.3. Plan for the article

The goal of this article is to put pressure on the harm footprint argument for moral veganism whether the purveyor of the argument takes position (A), (B), or (C). The argument will proceed in two stages, uncovering *the vegan's dilemma*.

First, I will argue that if the harm footprint argument (in any of the three guises) is successful in showing that producing and consuming animal products is morally impermissible, then parallel harm footprint arguments also succeed in showing that a wide range of practices is morally impermissible. But among these is a vast array of practices that make contemporary life uniquely valuable and that we do not think are morally impermissible. If the harm footprint argument for veganism is a good one, then it follows that vegans, by engaging in this wide range of modern practices, are living immorally *by their own lights*. That is the first horn of the vegan's dilemma.¹⁵

The first horn is an uncomfortable place to rest. Perhaps the harm footprint argument for veganism misfires. One might indeed think that the argument is reduced to absurdity by its implication that vegans and everyone else living a modern lifestyle are living atrociously immoral lives. The problematic assumption of the argument leading to this conclusion seems to be that morality requires not just taking harm *into account*, but *minimizing* harm. If we reject that assumption and assume, instead, that morality requires keeping the harm of our actions under some threshold or within some budget, then we can avoid the first horn of the dilemma but we land upon the second horn instead. On this second horn, we recognize that there are many ways of keeping our harm footprints within budget by making judicious selections from among the vast array of modern lifestyle practices available to us. Such practices, of course, include our dietary choices, but are not limited to them. On the second horn of the vegan's dilemma, therefore, veganism is not a requirement of morality.

2. Implications of harm minimization reasoning

In this section, I trace the implications of the harm minimization argument for veganism. I argue that if the harm minimization argument for veganism is a good one, then it follows that it is morally impermissible to engage in many common practices.

2.1. Travel

It is widely agreed that climate change is the most pressing environmental problem facing humanity. If it goes on unabated, climate change will produce massive, unnecessary harm to animals, the environment, and humans. It is also widely agreed that anthropogenic GHG emissions bear significant causal responsibility for climate change. According to the logic of the harm footprint argument, if we engage in practices that produce this massive harm when other suitable practices are available, then we behave

¹⁵This first horn of the dilemma may be seen as a sort of demandingness objection, so that part of the contribution of this article is to point out that this popular and influential argument that veganism is obligatory has to confront this well-known objection.

immorally. With respect to GHG emissions, one such practice is the consumption of meat from ruminants, such as cattle, sheep, and goats. Not only do they require feed, which is often grown with the use of GHG-emitting farm equipment, but ruminants emit methane, a GHG, as a by-product of enteric fermentation. So some conclude on the basis of harm minimization reasoning that producing and consuming lamb (the highest GHG-emitter), for instance, is morally impermissible just in light of the associated GHG emissions, for other, less harmful sources of protein are available.

Air travel also produces very large GHG emissions. For example, the GHG emissions from a round-trip flight from Chicago to London once per year are the same as eating about two quarter-pound (113 grams) lamb burgers every day for one year, on top of one's normal diet.¹⁶ Flying such a distance for travel to academic conferences is standard practice for many academics. Like eating an extra half-pound of lamb per day, such travel is completely unnecessary – with possible exceptions for job seekers and other early-career academics. There are suitable alternatives, such as conferencing through video and getting feedback on one's work from colleagues in the philosophy blogosphere or via e-mail. Thus, if one is required, by harm minimization reasoning, not to eat an extra half pound of lamb per day just in light of the associated large and harmful GHG emissions, by parallel reasoning, one is required not to engage in such unnecessary air travel.

What goes for travel to academic conferences goes a fortiori for GHG-emitting leisure travel. If we are required to minimize our harm footprints, then even a habit of traveling across town to visit a friend or to see a movie is impermissible, for friends can be Skyped and movies can be downloaded.

Notice that the argument against inessential travel goes through on all three positions about causal efficacy under consideration. If (A), our actions are causally efficacious in the production of animal harm from our dietary choices, it seems just as reasonable to assume they are efficacious in the production of GHGs. One's weight on a plane really does cause it to burn more fuel. If (B), our dietary choices and travel choices are (usually) causally inefficacious but the threshold chicken response succeeds, then there is a parallel threshold flyer response: My decision not to engage in inessential air travel probably will not cause the airlines to offer fewer flights and produce fewer GHG emissions. But there is some large number of flyers such that if that number of flyers stops inessential air travel, the airlines will offer fewer flights and therefore produce considerably lower emissions and less environmental harm. So even though I will almost certainly not make a difference, there is some small probability that I will make a big difference. So my refraining from flying is morally and mathematically equivalent to reducing my proportional share of the GHG emissions associated with the flights I refrain from taking. If (C), our dietary choices and travel choices are causally inefficacious but the complicity response succeeds for dietary choices, then there is a parallel complicity response in the case of inessential travel. Even though the massive and unnecessary harm produced by widespread inessential travel will be produced whether or not I engage in inessential travel, I am complicit in that harm when I participate, so it is wrong for me to do so.

One might respond, first, that none of what I have pointed out should be surprising, much less a *reductio ad absurdum* of harm minimization arguments. Many environmentalists are more than willing to bite this bullet, for they do believe that we need massive change to our lifestyles in order to curb climate change. Second, one might claim,

¹⁶Bruckner, *Small-Scale Animal Agriculture*, p. 206.

there is something unsatisfying about the argument based on GHG emissions above. Producing and consuming animal products with high GHG emissions causes not just the harm of those emissions, but *additional* harms to animals, the environment, and humans, some of which were outlined at the beginning of the article.

While I do not think that this second response is ultimately effective, it does point to a common mistake in harm minimization reasoning. Let me explain the mistake, and then explain why I have not made it. The mistake is to concentrate just on one type of harm or one source of harm. For example, if we concentrate just on minimizing GHG emissions, then we will miss other morally relevant considerations. After all, one way to reduce GHG emissions would be to enslave a segment of the human population to ride stationary bicycles in order to produce electricity. This would not be an acceptable solution because of the massive harm it would do to those humans. Similarly, if we concentrate only on minimizing environmental harm more broadly and restrict our consumption of animal products to those produced organically, then we will ignore morally relevant animal welfare considerations that are left unaddressed by (U.S.) organic regulations. If we concentrate just on animal welfare and go indiscriminately vegan, then we will ignore relevant environmental considerations about different plant agriculture practices. We will also ignore considerations of human harm, such as to migrant workers picking strawberries in California or the victims of drug cartels that dominate the avocado trade.¹⁷

Although the second response above points out this common mistake, my argument does not make it. It is true that lamb consumption is attended with more harm than just GHG emissions. But my argument does not depend on the assumption that GHG emissions are the only harms of lamb consumption. The comparison with lamb is only for dramatic effect to illustrate the magnitude of GHG gas emissions from flying, and does not claim that the *total* harm footprints of the round trip flight and the lamb burgers are equal. My argument is just an instance of harm minimization reasoning.

To see this, recall the argument form sketched in the introduction: Suppose we have a practice P1 that produces massive harm to animals, the environment, and humans. Suppose that the harm that comes from P1 is unnecessary in the sense that another practice P2, which is a suitable alternative to P1, is available and produces less total harm. Then the harm footprint argument concludes that P1 is morally impermissible. Here is the current instantiation of that argument form: Flying the equivalent of a round-trip from Chicago to London for academic conferences produces massive harm to the environment (and therefore humans and animals). The harm that comes from such flying is unnecessary in the sense that other available practices (video conferencing, etc.) are suitable alternatives to such flying and produce less total harm. The harm footprint argument concludes that such flying is morally impermissible.

2.2. Pets

Pet dogs and cats are fed animal products. Indeed, there are so many dogs and cats in the U.S. that they consume about 25% of all of the products of animal agriculture that

¹⁷Mark Budolfson is the first philosopher (of whom I am aware) to give serious consideration to the human worker harm associated with various vegan staples. See Mark Budolfson, Consumer Ethics, Harm Footprints, and the Empirical Dimensions of Food Choices, in *Philosophy Comes to Dinner*, ed. by Chignell, Cuneo, and Halteman, pp. 163–81 (pp. 164, 168, and 171–72) and Mark Budolfson, Food, the Environment, and Global Justice, in *The Oxford Handbook of Food Ethics*, ed. by Anne Barnhill, Mark Budolfson, and Tyler Doggett (Oxford: Oxford University Press, 2018), pp. 67–94 (pp. 88–89).

are produced in the U.S.¹⁸ So keeping dogs and cats as pets is responsible for about 25% of the total harm to animals, the environment, and humans from animal agriculture. So keeping pet dogs and cats is massively harmful. Even if these pets were fed a vegan diet (and it is questionable whether cats' nutritional needs can be met on a vegan diet),¹⁹ producing crops to feed pets still causes massive harm. It causes massive harm to the environment by using land, water, fertilizer, and fossil fuel. It causes harm to field animals by killing or maiming them in the process of field operations for cropping.²⁰ There are suitable alternatives to keeping dogs and cats as pets. People could keep smaller pets with much lower harm footprints, such as hamsters or mice. People could replace relationships with pets with relationships with humans that already exist, such as orphaned children in need of adoption. These are all suitable alternatives that are less harmful. Therefore, harm minimization reasoning implies that keeping dogs and cats as pets is morally impermissible.

Again, this argument is perfectly parallel to the harm footprint argument for veganism on all three positions on causal efficacy. If (A), human dietary choices are causally efficacious, then pet dietary choices are also and the harm footprint arguments in both cases are equally strong. It will not do to respond that some pet food is composed of by-products not consumed by humans, such as hearts, livers, testicles, kidneys, spleens, and tripe, so that abstaining from keeping pets would not reduce demand for agricultural animals, but only cause a problem of by-product disposal. For those by-products certainly could be, and probably would be, processed into products attractive for human consumption (think hot dogs, Spam, and scrapple) if there were not a demand for pet food. If (B), individual dietary choices are not causally efficacious but the threshold chicken response succeeds, then an individual choice whether to keep a pet will also be causally inefficacious in producing or preventing animal harm, but a parallel threshold pet-keeper response will succeed. If (C), individual dietary choices are causally inefficacious but the complicity response succeeds, then an individual choice whether to keep a pet will also be causally inefficacious, but a parallel complicity response will apply.

Again, some will claim that this should not be surprising and that we should of course give up unnecessary travel and keeping dogs and cats as pets. Morality is simply very demanding. It sometimes requires that we make what we consider to be great sacrifices. Indeed, one omnivorous response to the veganic harm minimization argument is that giving up meat would be a great sacrifice.²¹ A reasonable veganic rejoinder is that great sacrifices are sometimes required by morality. So as yet, we seem to have no real embarrassment to veganic harm minimization reasoning.

¹⁸Gregory S. Okin, Environmental Impact of Food Consumption by Dogs and Cats, *PLoS ONE*, 12 (2017), 1–14 (p. 7).

¹⁹Kathryn E. Michel, Unconventional Diets for Dogs and Cats, *Veterinary Clinics of North America: Small Animal Practice*, 36 (2006), 1269–81 (pp. 1275–77).

²⁰The work that started much of the conversation about harm to field animals among philosophers is Steven L. Davis, The Least Harm Principle May Require That Humans Consume a Diet Containing Large Herbivores, Not a Vegan Diet, *Journal of Agricultural and Environmental Ethics*, 16 (2003), 387–94. For the state of the discussion on field animal harm at present, see Bob Fischer and Andy Lamey, Field Deaths in Plant Agriculture, *Journal of Agricultural and Environmental Ethics*, 31 (2018), 409–28.

²¹One author who can be read as making this response is Loren Lomasky, Is it Wrong to Eat Animals?, *Social Philosophy and Policy*, 30 (2013), 177–200.

2.3. Alcohol

Harm minimization reasoning implies that producing and consuming alcoholic beverages is immoral. According to one well-researched source, the production of alcoholic beverages globally requires about 54 million acres of cropland just to grow the grains and grapes used in alcohol production.²² Just like cultivating crops to feed vegan pets, cultivating crops for alcohol production results in massive harm to the environment and to field animals. Alcohol also produces great harm to humans. According to the Centers for Disease Control and Prevention, about 107,000 people die in the U.S. each year due to the consumption of alcohol – either their own or someone else's.²³ While alcohol consumption prevents some deaths due to the positive effect of moderate use on cardiovascular health, an honest accounting shows that the net effect on humans of alcohol consumption is very negative.²⁴ So the production and consumption of alcoholic beverages result in massive harm to animals, the environment, and humans. The harm that comes from the production and consumption of alcohol is unnecessary in the sense that other, less harmful beverages are available as suitable alternatives to alcoholic beverages. The harm footprint argument concludes that the production and consumption of alcohol is morally impermissible.

One might think that this argument does not apply to someone who consumes alcoholic beverages in moderation. The moderate drinker is not harming his own liver, causing alcohol-related traffic injuries or fatalities, or harming others by abusing them while drunk. The moderate drinker is not causing any harm to animals, humans, and the environment, because only a tiny quantity of crops is needed to produce the alcoholic beverages he consumes, and those beverages would be produced whether he drank them or not.

Clearly this response is just a version of the causal inefficacy objection. To answer it, we just need to consider again the three cases. In case (A), the moderate drinker and the occasional carnivore are both causally efficacious. In that case, the harm footprint argument applies equally to moderate drinking and occasional carnivory. Case (B) admits that the vast majority of moderate drinkers produces no harm, but gives a threshold drinker response. There is some number of moderate drinkers, such that if that number stopped drinking, then the alcohol industry would produce less alcohol and demand fewer crops, thereby reducing harm to animals, humans, and the environment from raising those crops. So by stopping drinking, one has a small chance of preventing a great harm, which is morally and mathematically equivalent to preventing the amount of harm proportional to one's moderate consumption. Case (C) also admits causal inefficacy, at least for the vast majority of cases, but claims that drinking alcohol makes one complicit in the great harm caused by the alcohol industry just as one is complicit in the harms of animal agriculture even if one's occasional consumption of meat makes no difference.

The alcohol-lover may respond that abstinence from alcohol would come at a great sacrifice of the pleasures associated with its consumption. Again, though, the usual

²²Oliver Klimek, The Environmental Impact of Alcoholic Drinks Production, *Dramming: Everything Whiskey* (August 22, 2014), <<http://www.dramming.com/2014/08/22/the-environmental-impact-of-alcoholic-drinks-production/>> [accessed 19 December 2019].

²³Centers for Disease Control and Prevention, Alcohol-Related Disease Impact (ARDI) Application, (2013), <www.cdc.gov/ARDI> [accessed 19 December 2019].

²⁴For the accounting, see Donald W. Bruckner, Gun Control and Alcohol Policy, *Social Theory and Practice*, 44 (2018), 149–77 (pp. 164–65).

veganic argument is not refuted by the parallel response of the omnivore that abstinence from animal products would come at a great sacrifice of the pleasures associated with their consumption, so this response by the alcohol-lover does not succeed either.

The alcohol-lover may also point out that the harms that result from the production and consumption of animal products are surely far greater than the harms from alcohol production and consumption. Even if this is true, it does not defeat the harm minimization argument against alcohol any more than it would defeat the parallel argument against animal products to point out that some other practice is even more harmful than the production and consumption of animal products. What matters to the harm minimization argument is just that the harm of a practice is massive and unnecessary, and it is so in the case of alcohol.

2.4. Procreation

The harm minimization argument for veganism is concerned with anthropogenic harms to animals, the environment, and humans. Anthropogenic harm is caused, of course, by humans. Holding all else equal, the more humans there are, the more anthropogenic harm is produced. So to participate in the practice of procreation (even by producing just one human) is to participate in a practice that causes massive harm to animals, the environment, and humans. The practice is unnecessary for any given individual, since there are plenty of suitable and less harmful things one can do with one's life other than to raise children. Instead of spending large amounts of time and money raising children, one could spend time volunteering for any number of causes to reduce harm to animals, the environment, and humans and donate money to such causes. We need not worry that the human race will die out because too many people will be motivated to abstain from procreation, so it is not necessary for any given individual to engage in the practice of procreation in order for the race to survive. Thus, procreation is a practice that is massively harmful to animals, the environment, and humans and that is unnecessary. It follows by harm minimization reasoning that engaging in that practice is morally impermissible.

It appears that the causal inefficacy objection does not apply to procreation, since procreation does cause a human to exist, who does cause harm. If for some reason I do not see, the inefficacy objection does apply, then threshold human and complicity responses will apply.

2.5. Occupations

The professorial occupation is not very harmful, at least in comparison to many other occupations. Yet when one is employed as an academic, one fails to engage in a practice that is even less harmful. One could, instead, be employed by a social services agency that provides family planning and contraception in order to help prevent the births of so many humans that cause so much harm to animals, the environment, and humans. Alternately, one could switch to a much more lucrative occupation and maximize one's earnings in order to give as much as possible to effective charities. As in the case of human procreation, any one individual's engagement in the professorial practice is not necessary for academic culture, knowledge, and institutions of higher learning to survive. Like humans, we academics are overpopulated, at least in the humanities. So massive harm to animals, the environment, and humans is done by participating in the professorial occupation, and it is unnecessary for any given individual to engage

in it. By the harm minimization argument, engaging in the professorial practice is immoral.²⁵

2.6. *The vast array of harmful practices*

Beyond the specific harmful practices discussed above, our lives contain a vast array of practices that, taken together, produce massive harm that is unnecessary. Keeping any animals for companionship or recreation, especially ones with large appetites such as horses, is very harmful to animals of the field that are killed or maimed through producing food for the companion animals. It is harmful to the environment due to the environmental cost of the farming that produces their food. It withholds a benefit from humans, who could be fed with crops grown on the land currently used to graze horses or grow crops for them. So recreational horseback riding is morally impermissible, as is owning, breeding, training, and boarding horses.

Using air conditioning in the summer if our health does not require it and heating our homes and workplaces above 55 or 60 degrees in the winter rather than wearing more and warmer clothes are also massively harmful to the environment and unnecessary, and therefore morally impermissible. Even if the energy we use for heating and cooling is electricity generated from wind or solar energy, windmills and solar panels take land that could be used for wildlife habitat. Moreover, the windmills and solar panels are produced with scarce natural resources, such as minerals used to make steel.

It is also unnecessarily harmful to eat any food other than mostly raw freegan or vegan food that provides necessary sustenance but without unnecessary embellishments such as oil, sugar, salt, spices, and herbs, all of which require environmental resources to produce. By the same reasoning, eating dessert is forbidden by morality.

Using life-extending technologies for humans is also wrong. It produces harm because by living longer one makes things worse for humans by polluting and using scarce resources, and harms animals even through eating vegan fare.

The list could be continued. It should be clear enough at this point, however, that the logic of the harm minimization argument extends very broadly. It seems to imply that it is wrong to participate in the vast array of practices that make modern life distinctly valuable. The harm minimization argument applies whether (A) our actions are causally efficacious, or if they are not but (B) threshold arguments apply, or (C) complicity arguments apply. Therefore, a vegan proponent of the harm minimization argument for veganism who lives a modern lifestyle is living a terribly immoral life by the standards of that argument.

3. Moving beyond harm minimization

3.1. *Alternative conclusions*

This dialectic began with the harm minimization argument. This argument was supposed to be distinctive by proceeding from uncontroversial moral assumptions. On

²⁵Working as a professor (rather than in social services, for example) seems like a failure to benefit rather than a harm. Yet, some have argued that some failures to benefit are indeed harms (Neil Feit, *Harming by Failing to Benefit, Ethical Theory and Moral Practice*, 22 (2019), 809–23). Moreover, philosophers employing the harm minimization argument for veganism often treat failures to benefit as on a par with harms, for example, in failing to use agricultural land more efficiently by producing more food to benefit humans rather than food for animals that humans then consume. Space precludes exploring these and other complications.

DeGrazia's version, the assumption was that animals have at least some moral status, while on that of Hooley and Nobis, the assumption was that it is wrong to cause serious harms unless those harms are justified. The similar assumption I have used in the version here is just that it is wrong to support or engage in a massively harmful practice that is unnecessary. We have now seen that the vegan supporter of such arguments has to bite a lot of bullets. She has to bite the bullets that procreation, keeping pets, and drinking alcohol are wrong, and that she is engaged in systematic wrongdoing by engaging in a host of contemporary lifestyle practices. So although the proponent of the harm minimization argument can save it by biting these bullets, doing so amounts to giving up what was supposed to be distinctively interesting and persuasive about the harm footprint argument in the first place, that it is consistent with common-sense morality.²⁶ While the starting points of the argument are consistent with commonsense morality, what follows from them is not. So perhaps there is something wrong with the starting points after all.

Rather than biting the bullets that follow from harm minimization reasoning, an alternative conclusion to draw from the hardness of those bullets is that morality does not require harm minimization. One might think that perhaps morality is demanding, but that it cannot be so demanding that it requires us to jettison all of the practices discussed above that give modern life so much of its distinctive value. Yet the harm minimization argument is intuitively appealing, so we should investigate what the argument gets right and where it might go off the rails.

One plausible position is that we should take seriously the starting point of harm minimization arguments. The starting point is the claim that harm is morally relevant. Harm should be considered in moral decision-making. The harm minimization argument seems to go off the rails by requiring not just the *consideration* of harm but the *minimization* of harm. Even if we agree that harm is morally relevant, it is quite a leap to infer that one must cut out all the harm that one possibly can. Another position, the one I will investigate here, is that one should not cause harm in excess of a per-person threshold or budget.²⁷

3.2. The size of one's harm budget

Immediately the difficult question arises: What is an appropriate harm budget? Full treatment of this question would require another article. I will make just one claim about the appropriate size of the budget: If we suppose that creating one child is morally permissible, then an individual's harm budget is greater than the excess harm associated with being an omnivore rather than a vegan. The reason is that procreation creates far more harm than the additional harm associated with one individual having an omnivorous diet rather than a vegan diet.

To see this, notice first that if one engages in procreation, then one is not causally responsible for all of the harm wrought by the child created. Since procreation requires two people, each person is responsible for only 50% of the harm wrought by the child.

²⁶Thanks to Bob Fischer for this point and this way of putting it.

²⁷Mark Budolfson also rejects harm minimization arguments (though for different reasons) and introduces the idea of a harm budget. See Budolfson, *Consumer Ethics*, p. 167. Budolfson's work emphasizes the self-defeat of veganic harm minimization reasoning in support of a veganic *diet*, because a careful omnivorous diet has a lower harm footprint than typical vegan diet. By contrast, my emphasis is on the overgeneralization of harm minimization to other modern *lifestyle practices* more generally.

So if one procreates, one produces another 0.5 human's worth of harm. If one lives in a high-income nation such as the U.S., Canada, Australia, or the U.K., where birth rates are relatively low and the average woman produces 1.7 children,²⁸ then one's child can be expected to be responsible for half of 1.7 children, or 0.85 children. So after only two generations, one's procreation is responsible for half of one's child, plus half of one's child's half of her 1.7 children, which is $(0.5) + (0.5)(0.85) = 0.925$ of a human's worth of harm.²⁹

So even if we count only the first two generations and we assume that you have only one child, you are responsible for almost one full human's worth of harm, in addition to the harm you would have caused had you not procreated. It seems safe to say that one full human (or even 0.925 of a human) produces, on average, more harm than the marginal harm that one human produces by having an omnivorous diet rather than a vegan diet. If we make the plausible assumption that the procreation of one child is morally permissible, then procreation must not produce a morally impermissible amount of harm. That is, procreation does not cause one to exceed one's harm budget. Since less harm is produced by an omnivorous diet than by procreation, having an omnivorous diet does not cause one to exceed one's harm budget either.³⁰

One might object: This calculation allocates the harm caused by a child to that child's parents. A child, once grown, is an autonomous being. Generally, one is not morally responsible for the actions and consequences of actions taken by another autonomous being. So it is not justified to allocate the harm caused by a child to that child's parents, much less that child's children, etc.

Grant that an autonomous being is causally responsible for her own actions. She is the proximate cause of her own actions. Suppose, however, that one has a causal role in setting up the conditions that lead another autonomous being to be the proximate cause of some harm. Then in some circumstances and to a considerable extent, one is morally responsible for the harm, even though one was not the proximate cause of the harm. For example, if I outfit rebels with munitions and they make free choices to kill many innocents in their quest to overthrow the legitimate government, then I may bear some moral responsibility for the deaths of those innocents. I was not the proximate cause of those deaths, but I played an integral role in the causal chain. So, at least if I could foresee the harm the rebels would wreak, I bear some moral responsibility for the harm, even though I was not the proximate cause of the harm. Similarly, if I create an autonomous being whose existence is a net harm to animals, humans, and the environment, knowing that that being will freely choose to cause harm and will likely freely choose to make yet more such beings, then I bear some moral responsibility for the resulting harms, even though the proximate causes of all of those harms are autonomous beings other than me.

If this calculation and the rejoinder to the objection are successful, then something significant follows: If procreation is morally permissible, then so is omnivory. If these

²⁸See The World Bank, Fertility Rate, Total (Births per Woman), <<https://data.worldbank.org/indicator/SP.DYN.TFRT.IN>> [accessed 19 December 2019].

²⁹If we continue to the next generation, one will be 50% responsible for one's child's 50% responsibility for her child's 50% responsibility for 1.7 children, or $(0.5)(0.5)(0.85)$. A little mathematics shows that if we continue indefinitely, the total number of humans one is responsible for producing by having just one child is $(0.5) + (0.85)((0.5) + (0.5)^2 + (0.5)^3 + \dots) = (0.5) + (0.85)(1) = 1.35$ humans.

³⁰Thanks to Bob Fischer for helping me see this point about the size of one's harm budget and for raising the objection immediately following.

are not successful, then none of the argument to follow is affected, since the argument to follow does not depend on any assumptions about the size of one's harm budget, the relative harms of omnivory and procreation, or the allocation of harm from procreation.

3.3. Alternatives to veganism

One lesson of Section 2 is that food is not special. While the harms of our dietary choices are important, they are not any more important than the harms that come from the many other activities and practices that are central to our modern lifestyles. We are now considering the position that we are not required to minimize harm, but to keep it within budget. Now that we see that food is not special, we are freed from a misguided preoccupation with dietary choices as unique and free to consider ways of keeping our harm footprints within budget other than through dietary choices.

For illustration purposes, it might be useful to focus artificially on several dimensions along which individuals' harm footprints can vary. Let us include diet, number of children, travel choices, occupation, and pastime activities. Consider the following people.

- Person A:* a vegan with no children who travels frequently for a job as a humanities professor and who spends free time hiking with her dog.
- Person B:* an omnivore with two children who travels infrequently, works as a social worker providing low-cost contraception, and creates usable objects from garbage as a pastime.
- Person C:* an omnivore with no children who travels infrequently, works as an environmental engineer mitigating air pollution, and brews beer as a pastime.
- Person D:* a vegan with three children who travels infrequently, operates an organic vegetable farm that takes special care to avoid harming field animals, and has no pastimes other than children and farm work.

Person A does pretty well and is likely within a permissible harm footprint. B has two children, which is very harmful,³¹ but this harm is more than offset by the pregnancies prevented by B's occupation. B is likely within a permissible harm footprint as well. C's worst harm is caused by C's diet, but C also seems within budget because she has no children. Although D scores very well on most dimensions, D has three children, which is probably enough to put D's overall harm footprint over budget. A possible accounting for the harms on these dimensions and in totality is provided in Figure 1.

To be clear, I do not think these are all the dimensions that matter. This is just a set of toy examples to illustrate the central point that food is not where all the action is. As we saw in 2.1, we see again that it is a mistake to concentrate on just one type or source of harm. For the sake of argument, we can grant the assumption that an omnivorous diet is more harmful than a vegan diet. We can grant what follows from this, that *all else equal* an omnivore has a higher harm footprint than a vegan. From this, very little follows about the morality of the two lifestyles. Very little follows because a *particular*

³¹I said in 3.2 that my argument here does not depend on any assumptions about allocating the harm of procreation, but here I am allocating the harm of a person's child to the parent, which may not be warranted if my response to the objection there does not succeed. But I need not allocate *all* of the harm of a person's child to the parent here. The reasoning here will go through even if we only allocate to the parent the harm footprint of the child until the child reaches the age of majority, say.

omnivore and a *particular* vegan (i.e., when *not* all else is equal) may both be living permissibly within an acceptable harm budget.

		Diet	Children	Travel	Occupation	Pastime	Total Harm	Within Budget?
A	Descriptor	Vegan	Zero	Frequent	Professor	Dog	Low/ Med	Yes
	Harm Footprint	Low	Low	High	Med	Med		
B	Descriptor	Omnivore	Two	Infreq	Soc. Work	Crafts	Low/ Med	Yes
	Harm Footprint	High	High	Low	V. Low	Low		
C	Descriptor	Omnivore	Zero	Infreq	Env. Engr.	Zymurgy	Med	Yes
	Harm Footprint	High	Low	Low	Low	Med		
D	Descriptor	Vegan	Three	Infreq	Org. Farm	None	High	No
	Harm Footprint	Low	V. High	Low	Low	V. Low		

Figure 1. Harm footprints.

3.4. *Objection: food is special*

I have claimed that food is not special because the harms of food production and consumption practices are on a par with the harms of many other modern lifestyle practices. One might question this crucial move by claiming that food *is* special.

The objection begins by pointing out that there are two relevant dimensions to consider when judging departures from harm minimization. The first dimension is (A) the degree to which the chosen option departs from harm minimization, that is, how suboptimal it is. The second dimension is (B) how much sacrifice or effort would be required of the agent in order to minimize harm. If we apply (A) and (B) to the cases of fossil-fueled travel and child rearing, for example, we will see that someone who travels and has a child has a much better excuse for doing those harmful things than someone who consumes harmful animal products. For many people, forgoing fossil-fueled travel and procreation would represent a significant personal sacrifice (B) which, once accounted for, shows that engaging in those practices may be suboptimal, but not as suboptimal as my original exposition allowed (A). Similarly, even admitting the dubious claim that forgoing the consumption of animal products imposes a significant gustatory cost, the health and financial benefits of a vegan diet are many times the gustatory costs (B). So once we account for these costs and benefits to the agent, veganism is even more beneficial than I originally allowed on (A). Thus, engaging in fossil-fueled travel and child rearing are not as far from optimal harm minimization as I originally claimed, but veganism is even more optimal (harm minimizing) than I originally claimed. Due to this difference along dimension (A), there is a better excuse for engaging in fossil-fueled travel and child rearing than I originally claimed, and a less good excuse for consuming animal products. In this sense, food is special, contrary to what I have claimed. In sum, there is something relevantly different about food, compared with the other forms of behavior that I have discussed. There is, in some cases at least, a significant cost to the agent in foregoing those behaviors. In the case of veganism, there is no net cost to the agent, and a significant benefit.³²

³²An anonymous referee raised this objection, using many of these same words.

Note first that this objection relies on many substantive empirical assumptions about the costs and benefits (to an agent and more widely) of travel, procreation, and consumption of animal products. For example, if it is dubious that it would be a great sacrifice of gustatory pleasure for some to give up eating meat, it is equally dubious that giving up the sorts of travel that I targeted above would entail significant personal sacrifices. Nevertheless, let us grant these assumptions for argument's sake. If we do, then we get the conclusion that food is at least a little bit special, for the reasons just rehearsed. For all that, the main argument of the article is not damaged. Food is special, but not special enough to get the vegan out of the dilemma.

To see this, consider the dilemma again. On the first horn, the harm minimization argument for the moral requirement to abstain from animal products is a good one. On that horn, we are also required to abstain from many other common lifestyle practices. If the objector is right, abstaining from travel and procreation would come at a very high cost to the abstaining agent. So the requirement to abstain is weaker than I originally claimed and failures to abstain from travel and procreation are more excusable than failures to abstain from animal products. Yet – and this is the important point – abstaining from travel and procreation *is still required*, because it minimizes harm.

Travel and procreation aside, there are still many other far-from-optimal but very common lifestyle practices, abstinence from which does not entail a great sacrifice to the agent. Alcohol consumption falls into this category. Participating in high-tech consumer culture, eating dessert, and building campfires probably do also. It follows that abstinence from all of these is morally required and failure to abstain has very little excuse. These implications still seem absurd. So even if we accept the objector's point that some failures to minimize harm are more excusable than others, the harm-minimization reasoning at the basis of the veganic argument under discussion still casts its net too broadly.

To escape this absurdity, we can again turn to keeping our harm footprints within budget rather than minimizing harm. Here, again, the specifics in the calculations of the harm footprints of Persons A through D may need to change in order to accommodate the objector's points about personal sacrifice and excuses, but the overall point does not: A lifestyle that includes the consumption of animal products but that is sufficiently low in harm in other areas can be within one's harm budget. Food may be special, but it is not special enough for the harm of a non-vegan diet to swamp or trump the harms of one's other lifestyle practices. So abstaining from consuming animal products is not a moral requirement, which is the second horn of the vegan's dilemma.

3.5. Procreation revisited

I made two claims above that one might question: that procreation is morally permissible and that the addition of a human to the world's population is a net harm. While these claims might be true for a randomly chosen parent or child, one might object to the first claim: To the extent that we can expect that parents will pass on many of their values and forms of behavior to their children, it is not clear that it is morally permissible for a bigoted war-loving meat eater to procreate. One might object to the second claim: It is similarly not clear that it is a net harm for a compassionate, tolerant, and morally conscientious vegan to procreate.³³

³³Again, thanks to an anonymous referee for this objection.

Before addressing the objection, it is important to note, first, that the main conclusions of the article do not depend either on the assumption that procreation is morally permissible for everyone or on the assumption that the production of a child by a compassionate, tolerant, and morally conscientious vegan is a net harm. I stressed in the last paragraph of 3.2 that the argument that followed in 3.3 does not depend on the claims I made in 3.2 about the size of one's harm budget, the relative harms of omnivory and procreation, or the allocation of harm from procreation.

Nevertheless, consider further reasons to think that procreation is a net harm, even for a compassionate, tolerant, morally conscientious vegan. While, all else equal, a child who inherits those traits will produce less harm than one without those traits, such a child living in the modern world will have a harm footprint that is largely insensitive to whether he has those traits or not. Although a conscientious person living in the modern world will seek to minimize his harm footprint, he will still have a harm footprint from the use of scarce and non-renewable resources, such as petroleum, natural gas, precious metals, landfill space, and agricultural land. He will contribute to climate change or slow its mitigation by emitting carbon as well as other GHGs such as those released from landfills or associated with consumer products such as refrigerants. He will contribute to air and water pollution. He will contribute to traffic congestion and crowding in cities or sprawl in suburban housing plans. Although he is a vegan, his diet will be associated with harm to field animals through agriculture. Wildlife will be harmed through habitat loss if he lives in a housing plan or travels on roads the construction of which displaced wildlife or results in vehicular collisions with wildlife. He may produce benefits in other areas of his life, but considering the vast array of largely unavoidable harmful practices of modern life, it is plausible that his existence will be a net harm to animals, humans, and the environment.

Again, although his existence is plausibly a net harm, nothing I have argued relies on that claim. To remove any final doubt, consider another example like those of persons A–D that does not rely on any assumption about the harm or permissibility of procreation. Consider E, a vegan who has a pet dog, drinks one bottle of wine per week, lives in a housing plan, commutes five miles to work daily, and travels long distance by air once or twice per year. Consider F, an omnivore who has no pets, drinks no alcohol, lives in a city high-rise, walks to work, and travels by air infrequently. Holding all else equal, E has a higher harm footprint. If we make the plausible assumption that E is living permissibly, F is also. So an omnivorous lifestyle is not impermissible.

4. Conclusion: the vegan's dilemma

The proponent of the harm minimization argument for veganism is in a dilemma. The harm minimization argument for veganism is a good one or it is not. On the first horn of the dilemma, the argument is a good one. On this horn, parallel arguments go through for the moral impermissibility of a vast array of modern practices that make modern life distinctly valuable. In this case, a vegan who engages in any of these practices is acting immorally by the logic of her own argument for veganism.

On the other horn of the dilemma, the harm minimization argument for veganism is not a good one. I suggested that it fails by moving from the recognition of the moral relevance of harm to the moral requirement of harm minimization and thereby proving too much. In this case, in which harm is relevant but not required to be minimized, the vegan proponent of the relevance of harm will have to admit that there is very little that is special, morally speaking, about the harm of dietary choices when there are so many

other ways of reducing one's harm footprint. On this horn of the dilemma, therefore, a vegan diet is not a moral requirement.

Thus, the vegan's dilemma is that either she is acting immorally by her own lights, or veganism is not morally required after all.³⁴

³⁴Thanks to Bob Fischer, Doug Portmore, Travis Timmerman, and anonymous referees for helpful written feedback and to Lynne Dickson Bruckner and Andrew Fenton for discussion. Thanks also to the participants at the conference 'The Future of Protein' at the University of Ottawa in October 2018 for discussion.