



Individual Responsibility, Large-Scale Harms, and Radical Uncertainty

Rekha Nath¹

Received: 29 November 2019 / Accepted: 31 December 2020 / Published online: 13 May 2021
© The Author(s), under exclusive licence to Springer Nature B.V. part of Springer Nature 2021

Abstract

Some consequentialists argue that ordinary individuals are obligated to act in specific, concrete ways to address large-scale harms. For example, they argue that we should each refrain from meat-eating and avoid buying sweatshop-made clothing. The case they advance for such prescriptions can seem intuitive and compelling: by acting in those ways, a person might help prevent serious harms from being produced at little or no personal cost, and so one should act in those ways. But I argue that such reasoning often relies on an overly simplistic assessment of the costs and benefits of those prescriptions, one that misconstrues or neglects important issues. Indeed, a closer look at those costs and benefits reveals just how little we often know about a number of real-world matters that bear on the expected consequences of the prescribed individual actions. Our predicament is one of *radical uncertainty*: we currently lack a sound basis for concluding that the relevant actions are more likely to do good than to backfire. I distinguish this empirically grounded objection from others, which are not convincing—including one based on the mere conceivability of the actions in question backfiring and another based on the supposition that such actions cannot make a difference in addressing the massive problems at issue. The upshot is that, at least for now, consequentialist arguments for many specific individual actions aimed at addressing large-scale harms are inconclusive.

Keywords Consequentialism · Utilitarianism · Vegetarianism · Animal ethics · Consumer ethics · Global justice

1 Introduction

Through our everyday choices and actions, many ordinary individuals in well-off societies seem to be implicated in the perpetuation of widespread, serious suffering. By eating factory-farmed meat, I might help support an industry that each year

✉ Rekha Nath
rekha.nath@ua.edu

¹ University of Alabama, Box 870218, Tuscaloosa, AL 35487, USA

inflicts extreme cruelty on billions of chickens, pigs, and cows. My purchase of a new pair of sneakers may help sustain deplorable sweatshop conditions endured by workers in distant countries. By running my air-conditioner and taking long-haul flights, I might contribute to warming the earth to ever more dangerous levels.

I will use the term “large-scale harms” to refer to problems like factory farming, sweatshop labor, and global warming, which involve avoidable, significant suffering that is perpetuated by human activities. I take it to be reasonably uncontroversial that addressing these large-scale harms, as well as others like them, is a morally urgent task. It also seems clear that, given the scale and complexity of these problems, the optimal means of addressing them is through institutions. But until the requisite institutional steps are taken, what responsibilities, if any, are borne by ordinary individuals? In particular, given the ways we find ourselves connected to the perpetuation of large-scale harms through our everyday actions, are there specific, identifiable ways in which each of us should modify our behavior so as to avoid contributing to the persistence of great suffering?

Some employ consequentialist reasoning to support an affirmative response to that question. They defend concrete prescriptions for how individuals should respond to large-scale harms—arguing that serious problems, such as factory farming, call for particular individual actions, such as embracing a vegetarian diet. The case they advance for such prescriptions can seem intuitive and compelling: by acting in specific ways, such as refraining from meat-eating, a person might help prevent serious harm from being produced in the future at little or no cost, and so one ought to act in those ways. But I will argue that, in many cases, such reasoning relies on an overly simplistic assessment of the expected costs and benefits of those prescriptions, one that misconstrues or neglects important issues. A closer look at those costs and benefits reveals that in these cases we often lack sufficient evidence to justify specific actions by appeal to consequentialist reasoning.

For much of this paper, I focus on the consequentialist argument for vegetarianism. I do so because that particular argument has been especially well developed. However, my primary concern is more general. Many consequentialist arguments for analogous conclusions—arguments that individuals should do this or that specific thing to address large-scale harms—face the same sorts of problems that I raise for the consequentialist argument for vegetarianism. I use the latter argument as a proxy for many arguments in that class.

The paper proceeds as follows. I begin in Sect. 2 by explaining the consequentialist argument for vegetarianism. In Sects. 3 and 4, I develop my argument that we currently lack the evidence that would be required for that argument to be cogent. I raise a number of problems involving the possibility that embracing vegetarianism might *backfire*, that is, the possibility that an individual’s abstention from meat-eating will make animal suffering even worse than it would otherwise be. The point is not merely theoretical. That it is *conceivable* that one’s action should backfire does nothing whatsoever to weaken the consequentialist case for acting that way: most any action could conceivably lead to any number of outcomes, positive or negative. But the backfire scenarios I discuss are not merely conceivable, fantastical possibilities. On the contrary, though improbable in absolute terms, the available evidence indicates that they are no less improbable than the possibility of one’s vegetarianism

achieving its desired results. In short, our predicament is one of *radical uncertainty*. We do not yet know enough to conclude that the cost-benefit analysis of vegetarianism definitively supports the case for individuals embracing that dietary choice to combat factory farming. The same is true for many other consequentialist arguments for specific individual actions aimed at addressing large-scale harms. I indicate why in Sect. 5, where I also raise further problems for those arguments. Until we pay greater attention to the sorts of real-world complexities that I discuss, many such arguments will remain inconclusive.

2 The Consequentialist Case for Vegetarianism

It might seem doubtful that any one person's vegetarianism could help reduce the suffering of factory-farmed animals. Indeed, it hardly seems reasonable to expect the factory farming industry, which slaughters tens of billions of animals annually, to take notice of the lower demand triggered by a single consumer's abstention from meat-eating, let alone to adjust its future output accordingly. The first step of the consequentialist argument for vegetarianism, then, requires explaining how, against this worry about causal impotence, a person could make such a difference through her dietary choices. Proponents of this argument include Almassi (2011); Matheny (2002); Norcross (2004) and see also (2008); Rachels (2011; Singer (2009[1975]) and see also (1980). To do so, proponents of this argument offer an example along the following lines.¹ Suppose that a chicken producer will adjust its future output if demand shifts by a margin of 100,000 chickens in a given production cycle. My choice about whether to buy a chicken therefore stands to make a difference if my purchasing the chicken would cause the total number of chicken purchases (in the requisite time-frame) to increase by a multiple of 100,000, and, conversely, my not purchasing the chicken would result in the overall tally of chicken purchases falling just one chicken short of the margin needed to prompt the producer to increase future production.

Although the illustration just set out shows that a person's purchasing decision could make a difference, the odds of it doing so are very low. Assuming that all the different possibilities of how many chickens are sold in the given cycle are equally probable, the chance of my decision shifting demand by the exact margin needed to elicit an industry response is 1 in 100,000. This brings us to the next step of the consequentialist argument for vegetarianism, which is to show why the case for this dietary choice holds in spite of these low odds (Norcross 2004: 233; Singer 1980: 335–336). As proponents of the argument explain, the lower the odds of my vegetarianism making a difference, the greater the amount of good my vegetarianism is positioned to do. So, in the case at hand, in the highly unlikely event that I make a

¹ Norcross (2004: 232–233); Rachels (2011: 886–887); Singer (1980: 335–336). Kagan (2011: 121–128) too discusses this reasoning, though his aim is not to argue for vegetarianism.

difference by not buying a chicken, that decision will produce an extremely morally significant result: it will spare 100,000 chickens from enduring immense suffering in factory farms in the next production cycle. A straightforward calculation reveals that the expected payoff of my decision not to buy a chicken in this case would be the same as it would be if I had absolute certainty that one fewer chicken would be factory farmed as a result of my making that choice: $0.00001 * 100,000 \text{ chickens} = 1 * 1 \text{ chicken}$ (Norcross 2004: 233). So, if one should avoid meat when doing so would be guaranteed to make such a difference, then one should also avoid meat in the case of low odds but a commensurately higher payoff, since the expected utility of that choice would be the same.²

So far, we have seen how consequentialist reasoning might show why one should not *purchase* meat. That a person should additionally avoid *consuming* meat relies on two further considerations. First is the straightforward observation that a person's consumption choices might influence purchasing decisions made by others. Second, by not eating meat, one might inspire others to reduce their meat consumption (Singer 1980: 336–337).

The final step of the consequentialist argument for vegetarianism is to show that any costs entailed by this dietary choice are likely to be outweighed by the good it might do. Here a difficulty arises, which is that once we take into account various factors that were set aside in the hypothetical example just considered, we see that we do not have a clear sense of how much good one's vegetarianism might do. In that simplified example above, we assumed that upon reaching some threshold, a shift in demand of a certain magnitude would translate into a commensurate shift in future output. But, as consequentialist proponents of vegetarianism acknowledge, in the real world the reduced animal suffering one's vegetarianism might help produce is likely to be somewhat lower—i.e. the expected payoff of my refraining from consuming a single chicken will not be equivalent to one fewer chicken being factory farmed in the future (Almassi 2011: 403; Rachels 2011: 886). For one thing, reduced demand for a product can lead to a drop in its price, which in turn can cause demand to rise. For another, there might be a sizable gap between how much demand for a product actually changes and how much of a shift in demand is needed to elicit an industry reaction. One reason for this discrepancy is that a certain margin of reduced demand will be too negligible to affect the orders that restaurants and supermarkets put in with producers, and thus some abstentions from meat-eating will not even stand to make any difference to industry output. Because it is not easy to measure the impact of such factors, consequentialist proponents of vegetarianism do not attempt to put forth even a rough estimate of the real-world expected payoff of a person's vegetarianism.

² Almassi's (2011) consequentialist argument for vegetarianism differs somewhat from the others. For Almassi, the case for embracing vegetarianism is not about how much good a person can thereby expect to do in absolute terms but rather involves showing that, in comparative terms, one can reasonably expect to be better-positioned to reduce animal suffering by opting to be vegetarian than by not. But his argument is no less susceptible to the concerns that I go on to raise for the others' reasoning.

If we don't have even an approximate sense of the expected payoff of vegetarianism, do we nevertheless have good reason to believe that the gains it might produce would outweigh its costs? To answer this question, consider some of the costs at issue (Champene and Merrell 2008: 58; Regan 1980: 310–311). Becoming vegetarianism might mean giving up foods one enjoys, and it may take some effort—requiring, for instance, that one learns to prepare meatless meals and figure out which restaurants serve decent vegetarian fare. Given the low cost of meat, it might also be more expensive to be vegetarian. There is the additional worry that getting enough protein in one's diet without meat could prove a challenge. Furthermore, by giving up meat, one might miss out on valuable experiences such as enjoying turkey at a Thanksgiving family gathering. Finally, some raise the worry that scaling back meat production could harm the numerous workers who depend on the meat industry for their livelihoods.

In response, consequentialist proponents of vegetarianism explain how this characterization of vegetarianism as costly relies on questionable empirical suppositions. With respect to the cited efforts of becoming vegetarian, they contend, the burdens involved tend to be minimal at best, and, in any case, most of those burdens would obtain only for a short period of time as one transitions to that new diet (Norcross 2008: 8; Singer 2009[1975]: 179). And, they observe, vegetarian fare is, in fact, readily available and often cheaper than non-vegetarian food (Norcross 2008: 9). These authors also dismiss concerns about difficulties in meeting basic nutritional requirements on a vegetarian diet (Singer 2009[1975]: 171, 178–182). Indeed, an overwhelming body of evidence points to the significant health gains associated with cutting meat from one's diet. The worry about missing out on valuable experiences invites the reply that there hardly seems to be a dearth of valuable opportunities for social interaction that do not revolve around meat-eating—surely, a person can meaningfully bond with loved ones without eating turkey (Norcross 2008: 9). As for the posited harms to meat industry workers, consequentialists note that these jobs rank among the worst in society (Rachels 2011: 891–892; Singer 1980: 334). Moreover, because reduced demand for meat would presumably coincide with greater demand for vegetarian food, we should expect lower meat production to be accompanied by an emergence of different economic opportunities in the food industry rather than a net loss in jobs (Norcross 2008: 9–10).

As for the concern about taste, these authors are skeptical that the difference in gustatory pleasure derived from a vegetarian diet compared to a non-vegetarian one would amount to much (Singer 1980: 333). Of the cited costs, however, this one seems hardest to dispute: granting that it would be better for people to not care so much about gustatory pleasure, many do seem to derive great enjoyment from meat-eating (Regan 1980: 309–310).³ Still, even if vegetarianism would entail a not-insubstantial loss in gustatory pleasure for some, these authors take the case for

³ Take the case of an Italian homeless shelter whose residents declared that they would prefer to live on the streets than to remain in a shelter serving them only vegan fare (*The Telegraph* 2016).

that dietary choice to hold—at least, for most people in most situations.⁴ They take this case to hold even in the face of considerable uncertainty about how much good one’s vegetarianism can be expected to do. This is because despite various factors that might weaken the influence one’s vegetarianism can have on industry output, it nevertheless seems reasonable to assume that one stands at least some chance of making such a difference by embracing that dietary choice (Almassi 2011: 403).⁵ And given the horrific nature of the animal suffering at issue, the reasoning continues, reducing such suffering even by a small margin would carry sufficient weight to overshadow the comparatively trivial costs entailed by deprivations in gustatory pleasure (Almassi 2011: 403–404; Norcross 2008: 7–10; Singer 2009[1975]: 171, 178–183 and 1980: 333–334). This suggests that consequentialist proponents of vegetarianism have a relatively light burden of proof to meet: so long as that dietary choice stands even a very small chance of bringing about a meaningful reduction in the misery endured by factory-farmed animals, the case for refraining from meat-eating holds.

3 Radical Uncertainty and the Possibility of Harm

For the most part, debate about the consequentialist basis for vegetarianism has focused on whether that dietary choice might do some good rather than standing no chance of making a positive difference.⁶ In this section, I turn to a further possibility, which has received comparatively little attention: the possibility that a person might do harm by embracing vegetarianism. Consider two illustrations of how that could happen.⁷

⁴ For discussion of relatively rare exceptions—for instance, involving people who cannot maintain good health without meat—see Almassi (2011: 403, 406–408); Norcross (2012: 384–388, 385–386).

⁵ Almassi (2011) discusses a different way one’s vegetarianism might do good: by avoiding meat, I might make dietary choices that help support better food-production systems like the plant-based meat industry. Thus my vegetarianism might help further reduce animal suffering since a thriving plant-based meat industry would seem to increase the likelihood of more consumers buying its products over meat. Although my main focus will be on the difference one might make by withholding support for harmful practices rather than providing support for beneficial ones, I illustrate how similar concerns arise in both cases.

⁶ For in-depth analysis of this matter, see Hudson (1993), Shafer-Landau (1994), Nobis (2002), and Chartier (2006).

⁷ In much of the ensuing discussion, I focus on the impact an American individual’s vegetarianism might have on domestic meat production. Not all consequentialist proponents of vegetarianism frame the issue this way. But doing so enables me to discuss salient country-specific factors that complicate an analysis of the real-world expected consequences of one’s vegetarianism. Along the way, I consider the extent to which points concerning the American context generalize to other contexts. Where such generalizations are not warranted, my arguments are accordingly limited in their reach. Still, given the sheer scale of American factory farming operations, even arguments with this narrower scope have significant implications. By recent estimates, around one in five of the approximately 50 billion animals factory farmed worldwide each year are reared in the U.S. And the overwhelming majority of land animals raised for food in the U.S., a staggering 99%, are factory farmed—a proportion that is unmatched by any other country, at least to date. Toward the end of this section, I take up the broader issue of how one’s vegetarianism might help reduce animal suffering on a global level.

Backfire 1. Adam, an American, refrains from eating meat, and this causes demand for meat to drop by the exact margin needed for the industry to take notice. However, rather than responding to this shift in demand by reducing its output in the next production cycle, a major meat producer in the U.S. seeks out new markets abroad in which to unload this unanticipated domestic surplus. The search is successful. Over the coming months, demand for American meat in foreign countries grows rapidly, prompting the industry to significantly scale up its output intended for export. Due to this expansion, the meat industry's net output grows at an even faster rate than it would have had domestic demand remained stable.

Backfire 2. As in the previous case, Adam's vegetarianism helps reduce demand by the exact margin needed to trigger an industry reaction. And this time, the reduction in demand leads to the desired outcome: it causes some factory farms to scale back output in the next production cycle, resulting in fewer animals enduring cruel conditions. Alas, this victory is short-lived. In response to publicity about the meat industry's decline—replete with doomsday predictions of mass lay-offs—a contingent of patriotic Americans fear that their nation's traditional bacon-double-cheeseburger-loving way of life is under siege by a radical anti-meat movement. These folks band together, resolving to do whatever they can to protect their cherished way of life. The social campaign they embark upon to promote greater meat consumption ("Eat meat at every meal if you love your country!") pays off, resulting in an unprecedented boom in meat production.

Now, that it is merely possible for one's vegetarianism to do harm does not undermine the consequentialist argument for that dietary choice. The relevant question is whether the risk of one's vegetarianism backfiring is great enough to outweigh the potential good it stands to do. It might seem that we are not currently in a position to answer this question. As we saw in the previous section, at present, there seems to be much we do not know about the expected consequences of a person's vegetarianism. We do not have the information we would need to work out with any precision the odds of that choice doing good, the odds of its backfiring, or the magnitude of good and harm associated with the various possible outcomes that it might help produce.

However, facing uncertainty about such matters need not in general stand in the way of a person's being able to reasonably conclude that a certain action is justified in light of its expected consequences.⁸ Consider an example.

Crosswalk. I am in a position to help a frail old woman cross the street at a busy intersection. It is very cold. Without my help, she may hesitate for a very long time at the curb, afraid to cross, and thus risk prolonged exposure to frigid temperatures. Moreover, in attempting to cross on her own, she may fall

⁸ One might disagree, regarding such uncertainty as a pervasive problem that stands in the way of our almost ever being able to reasonably conclude that particular actions are justified based on their expected consequences (Feldman 2006; Lenman 2000). But I set that position aside, since it is plainly incompatible with the consequentialist argument for vegetarianism.

and suffer serious injury. However, she might well make it across the street just fine on her own. It is also possible that my attempt to help her will backfire—for instance, as I am helping her, a drunk driver could run a red light, striking and killing her. I am unable to determine with precision the odds of these (and other) possible outcomes. And, unfortunately, the old woman is unable to communicate her preferences on the matter to me.

The fact that I cannot determine even the approximate odds of the various possible outcomes that might be produced by my attempting to help the old woman cross the street does not constitute a good reason against doing so, let alone a decisive reason. Background knowledge might adequately justify the judgement that I am more likely to do good than harm by assisting her. For example, I might know that it is an extremely rare occurrence in this area for drivers to run traffic lights and hit pedestrians. And I also might know that people as old as she appears to be are especially susceptible to suffering frostbite as well as to falling and sustaining serious injuries. I do not know the frequency with which such events occur, but such precision is not needed. By appeal to such considerations, I could reasonably conclude that I would be justified in assisting the old woman on the grounds that the relatively high chance of my doing good outweighs the very low risk of my causing harm.

But consider a different case:

Virus. You are extremely ill and are experiencing symptoms associated with a deadly novel coronavirus to which you might have been exposed. You do not have access to reliable testing to confirm whether you have this virus or are suffering from something else; a nasty strain of the flu has been going around, and such a diagnosis (which also cannot be confirmed by available testing) could account for your symptoms. You do, however, have the opportunity to try a new experimental drug. This drug has been taken by some patients thought to have this coronavirus but has not yet been tested in rigorous clinical trials. Initial evidence suggests that the drug may prove life-saving for patients battling this virus who would have had little chance of survival had they not taken the drug. But, initial evidence also suggests that taking it poses grave risks. In particular, it may be lethal when taken by patients who do not have the virus. Unfortunately, you do not have time to seek out further evidence. You must decide right away whether to take the drug.

In this case, you do not have sufficient information to decide how to act based solely on an assessment of the expected consequences. Whether you should take the drug seems to be a coin toss: available evidence does not provide a sound basis for expecting things to go better by acting one way over the other. The situation involves what I call *radical uncertainty*. In cases of radical uncertainty, a person's epistemic position concerning the odds of different possible outcomes being produced, or the amount of good or harm associated with different possible outcomes of an action under consideration, is so impoverished that one cannot reasonably conclude whether to perform the action based on its expected consequences. Radical uncertainty can be distinguished from *ordinary uncertainty*. *Crosswalk* involves ordinary uncertainty. In that case, despite my inability to determine even the approximate

probabilities associated with the different outcomes that might be produced by my helping (or not helping) the elderly woman, or to work out just how much good or harm those outcomes would entail, I have sufficient evidence to draw a reasonable conclusion concerning how I should act.

I will argue that we currently face radical uncertainty in the case of vegetarianism. That is, I will argue that based on what we know at present, we cannot reasonably conclude that an individual is more likely to do good rather than harm overall by embracing vegetarianism. It might seem otherwise. It might seem that this case involves merely ordinary uncertainty. Just as we may reasonably judge the risk of harm in *Crosswalk* to be very low, so too it might seem highly improbable that a person's vegetarianism would lead to harms of the kind that occur in *Backfire 1* and *2*. Let us examine this claim.

In *Crosswalk*, the risk of doing harm does not tell against helping the elderly woman all things considered because I have sufficient, albeit incomplete and imprecise, information to justify that action. I have good reason to believe that the odds of doing harm by attempting to help her are very low. Accordingly, the risk of doing harm appears to be far outweighed by the comparatively high odds of my doing good. But to dismiss the risk of doing harm in *Backfire 1* and *2*, it will not do to merely posit a parallel claim, that the odds of a person's vegetarianism leading to the sorts of harmful outcomes described in those scenarios are very low in absolute terms. Granting that there is a very low chance of one's vegetarianism backfiring in those ways, the risk of those (improbable) harms could still threaten the consequentialist case for vegetarianism. That is because consequentialist proponents of vegetarianism accept that the chances of one's vegetarianism doing good (rather than making no difference) also tend to be very low in absolute terms. To recall, they emphasize that such low odds do not weaken the case for vegetarianism, due to the immense gains one's vegetarianism might produce in the unlikely event that it makes a difference. But parallel reasoning holds true of the harmful outcomes in *Backfire 1* and *2*. In those cases, a person's vegetarianism leads to harms of a comparable magnitude to the potential benefits that consequentialist proponents of that dietary choice emphasize, namely, a great increase in animal suffering. Even if it is very unlikely that one's vegetarianism would cause such harms, given the magnitude of suffering they involve, the logic of the consequentialist argument for vegetarianism implies that a person deciding whether to embrace vegetarianism should take seriously the risk of producing them. For the risk of vegetarianism backfiring in those ways to be dismissed, it would need to be shown that in the unlikely event of one's vegetarianism making any difference at all, the chances of it doing good can be reasonably judged to be greater than the chances of it doing harm of a comparable magnitude.

Are we justified in concluding that a person is more likely to do good rather than harm by embracing vegetarianism? One might argue that we are, reasoning as follows. It might be thought that the consequentialist's explanation of how vegetarians can help reduce animal suffering is grounded in sound reasoning about how markets work. By contrast, the sorts of scenarios described in *Backfire 1* and *2* might seem to present nothing more than merely fantastical possibilities. That those harmful outcomes stand any meaningful chance of being produced by a person's vegetarianism

is not substantiated by empirical evidence. As such, those hypothetical scenarios do not weaken the case for that dietary choice.

Consider how this line might be developed in response to *Backfire 1*. In that case, Adam's vegetarianism helps produce a decrease in domestic demand for meat that, perversely, precipitates an increase in global demand, resulting in a net increase in animal suffering. One might question how realistic this scenario is on the following grounds. If the industry could increase its profits by expanding its overseas market share, then it would do so independently of a sudden fall in demand. The suggestion that the industry would be prompted to explore this avenue for gain specifically in response to a reduction in domestic demand is implausible. So, even if Adam's vegetarianism could in theory produce the harmful effect in question, it is highly doubtful that it would actually do so.

However, that reasoning is predicated on the assumption that in general economic agents tend to explore viable avenues for increasing their profits. But that assumption is doubtful. It is not difficult to explain why, in the ordinary course of events, an industry might not attempt to increase its profits by expanding its operations. Perhaps industry executives are reluctant to undertake the significant front-end costs and risks entailed by such an expansion—requiring, for instance, the construction of new facilities and the potentially lengthy ordeal of forging ties with foreign business partners. However, a drop in domestic demand could understandably provide the impetus for exploring this option. In fact, we have evidence revealing that this seems to be precisely what occurred with the tobacco industry in the past fifty years (Bolyky 2014; *The New York Times* 2008). In response to a downward trend in smoking in affluent nations sparked by anti-tobacco public health campaigns, that industry adopted aggressive marketing tactics in poorer nations. These efforts paid off. As tobacco use has decreased steadily in affluent nations, it has risen dramatically in many populous, lower-income nations. Tobacco companies could just as easily have made those efforts *before* smoking declined in affluent nations, and there is no particular reason to suspect the results would have been any different. But they did not. So, based on what has been witnessed in this parallel case, the harmful outcome of Adam's vegetarianism described in *Backfire 1* hardly seems so improbable, as compared to the probability of his vegetarianism doing good.

Similar points apply to a parallel worry that might be raised about *Backfire 2*, in which the short-lived reduction in industry output brought about by Adam's boycott turns out to be a pyrrhic victory due to the success of the pro-meat counter-movement that it spawns.⁹ Consequentialist proponents of vegetarianism observe that an individual's abstention from meat-eating might encourage others to act similarly. But they fail to cite evidence supporting the claim that a person's vegetarianism is more likely to get others to follow suit rather than to have the opposite effect.

⁹ Here for ease of language in drawing comparisons with like cases, I refer to a person's abstention from meat-eating as an instance of a *boycott*. To clarify: I use this term to refer to instances in which individuals deliberately withhold financial support from an industry through their consumption decisions and do so without necessarily coordinating with others acting similarly. Often, in fact, deployment of this tactic is accompanied by various forms of social activism, as well as by coordination among supporters of a cause. But my use of the term should be understood in the narrower sense specified.

Further, given the deep and steadily rising polarization of the American populace, it does not seem out of the question, or even especially unlikely, that a significant increase in the numbers of people embracing vegetarianism would lead to greater divide on this issue (Foran 2017).¹⁰ Consider a different example of a boycott backfiring in this sort of way. In 2012, sparked by outrage at the views expressed against same-sex marriage by the American fast-food chain Chick fil-A's CEO, numerous individuals staged a nationwide boycott of the restaurant. In response to this boycott, a counter-boycott was initiated, and this led to the company enjoying record high sales (Glickman 2012).¹¹ And another American fast-food chain, Arby's, has responded to other restaurants' adoption of plant-based products that imitate meat by seeking out innovative ways to use meat in the place of traditionally vegetarian offerings. Thus was born the "marrot," a meat-based food meant to resemble a carrot (Yaffe-Bellany 2019).¹² Additionally, there have been documented instances of the American meat industry using scare tactics to stir up public concern about a so-called "vegan conspiracy," in which those concerned about animal suffering are cast as extremists with a radical agenda (Norwood and Lusk 2011: 58). The industry's use of such tactics in the face of growing public opposition to factory-farmed meat might well help fuel a counter-movement of the sort described in *Backfire 2*. At least, that outcome would appear to be no *more* improbable than the positive outcome that vegetarians such as Adam hope to achieve.

So, by appeal to empirical evidence about like cases, we can conclude that the scenarios described in *Backfire 1* and *2* should not be dismissed as merely fantastical possibilities. But even granting that it is not completely out of the question that the harmful outcomes in those cases could come about, it might be thought that the potential good a person's vegetarianism stands to do clearly outweighs the harm it might produce. One difficulty involved in judging whether this is so is that (as we saw in the previous section), consequentialist proponents of vegetarianism do not provide us with even an approximate sense of how much good a person's vegetarianism can be expected to do. They do show that it is possible for one's vegetarianism to make a positive difference. Demonstrating that much suffices as a response to the causal impotence worry. However, responding to the backfiring concern calls for a weightier burden of proof. An adequate response to the backfiring concern requires showing that the expected good of one's vegetarianism is great enough to outweigh the risk of inadvertently doing harm by making that choice. In the face of concern about a given action backfiring, the observation that the action might do some good would not, by itself, warrant concluding that, all things considered, one could reasonably expect to produce better consequences by performing, rather than not performing, the action.

¹⁰ This concern may not apply to societies in which polarization is much lower than in the U.S. Still, even in less polarized societies, examining the prospects of this particular social issue being divisive seems worthwhile. Consider, for instance, evidence of vegans in the U.K. being widely perceived in a negative light (Cole and Morgan 2011; Reynolds 2019).

¹¹ Friedman (1999: 20) discusses another instance of a boycott, also involving a politicized issue, backfiring in this way.

¹² I am not making this up.

Without providing an estimate of the expected good of a person's vegetarianism, the consequentialist proponent of vegetarianism might argue that the backfiring concern does not undermine the case for that dietary choice by emphasizing that they have a fundamental principle of economics on their side: the well-established principle governing the relationship between supply and demand in free markets. On this principle, we generally have good reason to expect an industry to adjust its production levels to correspond to sufficiently sizable shifts in consumer demand. Thus, given the typical operation of markets, if one's vegetarianism helps reduce demand by a great enough margin for the meat industry to take notice and adjust its production levels, it would seem far more likely for that dietary choice to help trigger a net decrease rather than a net increase in industry output.

But is that expectation justified? This question is not adequately resolved by abstract reasoning about how markets ideally operate. Other factors are relevant. One is the efficacy of boycotts—in particular, whether boycotts, in general, tend to succeed in bringing about net reductions in production levels of the goods that they target. Of course, to conclude that they do, it would not do to merely cite some instances of successful boycotts. The existence of such cases is consistent with an overall trend of boycotting proving an ineffective, or even a counterproductive, tactic. And indeed, such a pessimistic assessment of the tactic may be warranted. One historian takes it to be “indisputably the case that the vast majority of boycotts throughout American history have been putative failures” (Glickman 2009: 2). Another scholar studying boycotts emphasizes the paucity of rigorous data of the sort that would be needed to draw conclusions about the overall effectiveness of boycotting (Friedman 1999). One specific challenge here is that in cases of apparently successful boycotts, it can be hard to determine the extent to which the gains at issue are attributable to boycotting (i.e. to a withdrawal of consumer support) independently of other tactics deployed by a social movement. Moreover, historical data on the matter may be skewed by bias. Plausibly, countless instances of failed boycotts have been long forgotten, while success stories are more likely to be remembered.

Further, in the case of the American meat industry, the prospects of a boycott's success might be diminished by the significant political clout and generous government support that this industry has long enjoyed.¹³ In recent years, there have been numerous instances of the government helping the meat industry through difficult times in which its production has exceeded consumer demand. In 2011, in an effort to help “bring supply in line with demand,” the U.S. government responded to a glut of chicken on the market by purchasing \$40 million worth of the industry's products, on top of its typical annual purchase from the industry of about \$100 million (Tomson 2011). Then, the following year, the American government again assisted the industry when it faced severe drought conditions, this time purchasing about \$170 million worth of meat and fish products (Doering 2012). Such massive government support provides a plausible explanation for how the meat industry's output could

¹³ The American context hardly seems unique in this respect. Meat production is heavily subsidized in many parts of the world in which factory farming is practiced, including in Australia, Brazil, Canada, the European Union, and the U.K. to name but a few such examples.

remain high even in the face of sizable reductions in demand. Thus, this is another factor, in addition to those mentioned in the previous section, that could well diminish the potential returns of a person's vegetarianism. Indeed, in light of this factor, the good one's vegetarianism could be expected to do might be only a small fraction of what would be the case in the absence of such government support. This is relevant because the lower its expected good, the more seriously the risk of harm threatens to undermine the consequentialist case for vegetarianism: one may not be justified in risking even a small chance of doing great harm if the expected good one might do by acting in the given way is extremely low.

The strong government support enjoyed by the American meat industry provides an additional reason for not dismissing the risk of one's vegetarianism doing harm in the way that occurs in *Backfire 1*. When the government purchases massive quantities of surplus meat from domestic producers, it must find ways to use it. Although it donates some to food banks, it has a strong incentive to recoup some of those significant expenditures. And one plausible avenue for doing so would be to find foreign markets on which to offload this product. The U.S. government has considerable bargaining power, thus making it better positioned to secure new avenues for exporting factory-farmed meat than the industry is (perhaps by leaning heavily on trade partners to import some such surplus in exchange for favorable terms of trade concerning other matters). Government efforts of this sort played a clear role in helping the American tobacco industry find new markets for its products in the face of falling domestic demand. For decades, the U.S. government actively negotiated trade deals that facilitated the tobacco industry's expansion into developing country markets (Bollyky 2014). These observations provide further support for why it would not be so improbable for the outcome in *Backfire 1* to come about, as compared to the probability of a positive outcome—thus lending support to the claim that the decision to embrace vegetarianism involves uncertainty of the sort that is present in *Virus*.

Suppose proponents of the consequentialist argument for vegetarianism concede that the possibility of one's vegetarianism backfiring in the described ways doesn't seem any more remote than the possibility of one's vegetarianism doing good. They might nonetheless maintain that their case remains on solid footing, reasoning as follows. If one's vegetarianism did backfire in those ways, the resulting harms should be regarded as merely temporary setbacks, which need not impede long-term progress. Imagine, for instance, that the harm of *Backfire 1* comes about: Adam's vegetarianism facilitates the emergence of a new export market for meat, thereby causing a net increase in production. Surely, that should not prompt him or others to abandon the vegetarian cause. Quite the opposite. We should instead react by trying to raise greater global awareness of the horrors of factory farming, perhaps focusing especially on societies in which factory-farmed meat is not (yet) consumed widely. Parallel reasoning applies to *Backfire 2*. If a counter-movement hostile to vegetarianism emerges, we should try to engage productively with opponents of vegetarianism and work to gain broader support for the movement in the future. Overall, then, such backfiring scenarios turn out to involve relatively minor, short-term harms when weighed against the much greater good that individuals' vegetarianism can do in the long term by helping to spare vast numbers of animals worldwide from torture.

But that response misses the point. All it does is describe hypothetical scenarios illustrating how it is possible for a person's vegetarianism to do greater good than harm in the long run, even if it initially backfires. An adequate response to the backfiring concern would have to show not only that such an envisioned outcome is *possible*, but moreover that it is *probable* relative to the alternative scenario in which one's vegetarianism does harm that is *not* offset by such future gains. As for the suggestion that efforts could be made to address inadvertent harms were they to arise, again we must consider the likelihood of such efforts succeeding and the likelihood of such efforts failing. That is an empirical matter, and there is no clear reason to conclude that it will resolve in favor of the former (successful) outcome. Let us briefly return to the tobacco case. Once the dangers of smoking were well established, reducing tobacco use worldwide became a key global public health imperative. However, as we now know, anti-smoking campaigns embarked upon in affluent societies starting in the 1960s had the unintended but devastating consequence of causing an enormous net increase in tobacco use globally, at the grave cost of tens of millions of deaths in the years since. Just as it would have been possible for that harm to have been mitigated once it became clear that tobacco use in developing countries was rapidly rising as developed-country use declined, well-intended public health efforts to do so have tragically seen little success (in no small part due to the tobacco industry's formidable political influence). Reflection on this case should serve as a cautionary tale against embracing optimistic projections in the parallel case of the meat industry, without careful consideration of relevant evidence that bears on their soundness.

I have argued that radical uncertainty undermines the conclusion that an individual should embrace vegetarianism based on the good that might do. The problem is not one of ordinary uncertainty, as in *Crosswalk*. Rather, the case is analogous to *Virus* insofar as one knows that by undertaking the action under consideration, one has some chance of producing great good but also some chance of doing serious harm. And as with *Virus*, in the case of vegetarianism, we simply do not have the evidence needed to support the judgment that acting in the given way is more likely to result in good rather than harm overall.

Let us clarify just how the situation of the individual contemplating vegetarianism is parallel to *Virus*, rather than *Crosswalk*, in this key respect. The relevant uncertainty does not concern how likely one's vegetarianism is to make a difference in addressing animal suffering. As consequentialist proponents of vegetarianism grant, the odds of one's vegetarianism making any difference are very low. The relevant uncertainty concerns whether one can reasonably expect her vegetarianism to be *more* likely to decrease rather than to increase animal suffering, in the unlikely event that it makes any difference at all. And we simply do not know whether that comparative proposition is true or false. This is just as in *Virus*. In *Virus*, the available evidence does not provide a sound basis for expecting things to go better (achieving a positive health outcome) by acting one way (taking the drug) over the other (not taking the drug). Likewise, in the case of vegetarianism, the available evidence does not provide a sound basis for expecting things to go better (decreasing animal suffering) by acting one way (embracing vegetarianism) over the other (not embracing vegetarianism).

On the one hand, appeal to the basic logic of how markets work explains how a drop in demand might cause future production levels to fall. On the other hand, because the economy does not operate in a vacuum, the principle governing the relationship between supply and demand is not the only consideration bearing on how future output might be affected by lower demand. On the basis of other well-known factors—e.g. government support of an industry; the politics of international trade; and the backlash social movements can invite, especially in a polarized climate—a consumer boycott of the meat industry could have exactly the opposite of its intended effect, leading to a net increase rather than a net decrease in meat production. The “could” here does not refer to a remote, merely theoretical possibility. Instead, in light of those factors, for all we know, the odds of the boycott backfiring may be roughly the same as those of the boycott succeeding. That one’s vegetarianism could, in that sense, be a part of a causal chain of events that results in a net increase in animal suffering is grounded in established empirical evidence about like cases in which such backfiring appears to have occurred. Moreover, the prospects of one’s vegetarianism backfiring in the described ways are grounded in empirical data that are about as well-supported as available evidence that can be advanced in support of that dietary choice doing good. The problem is that, just as in *Virus*, we lack sufficient evidence about the comparative likelihood of these two (unlikely) outcomes. We lack the sort of evidence that would be needed to conclude that the risk of harm is outweighed by the potential good one might do by embracing vegetarianism.

The foregoing reflections bring into focus how the individual contemplating vegetarianism faces a situation relevantly similar to that in *Virus*, and relevantly different from that in *Crosswalk*. We are able to point to factors that, on the one hand, explain how one’s vegetarianism might produce significant good. And on the other hand, we can point to factors explaining how one’s vegetarianism might result in serious harm. That is, we know that sometimes boycotts lead to a sustained drop in production of the boycotted goods, and we know that sometimes on the basis of the aforementioned factors boycotts backfire. Also like *Virus*, and unlike *Crosswalk*, we do not have a firm empirical basis—say, by appeal to general trends about cases relevantly similar to vegetarianism—to conclude that the action under consideration is more likely to do good than harm.

I do not claim that radical uncertainty poses an insurmountable obstacle to the consequentialist argument for vegetarianism. The relevant information is unknown but not unknowable. However, the argument will remain unconvincing unless the challenge I have raised here is answered. In particular, for such an argument to succeed, more serious attention must be given to the economic, social, and political complexities of the real world that I have discussed—complexities that imply that the possibility of one’s vegetarianism backfiring cannot be dismissed as fantastical.

4 Defining the Good

In this section, I will explain a further challenge to the consequentialist argument for vegetarianism. I will demonstrate how a person’s vegetarianism might not be justified *even* if it helps bring about a sustained decrease in the meat industry’s output

(and thus would avoid leading to harms of the sort produced in *Backfire 1* and 2). This is because the matter of how one's vegetarianism might do good is not as straightforward as we might initially take it to be.

Consider the following case.

Backfire 3. As in *Backfire 1* and 2, Adam's refusal to buy meat results in demand falling by just the margin needed for the industry to take notice. This time, unbeknownst to the public, the nation's largest meat producers had been on the verge of implementing significant changes to their operations geared at treating animals more humanely. However, owing to this unanticipated drop in demand, the industry's profits are much lower than expected. Consequently, industry executives decide that it is no longer financially viable to proceed with this expensive undertaking. In the years that follow, due to the reduction in demand that Adam crucially helped bring about, fewer animals are reared in factory farms than otherwise would have been, but the animals endure much greater suffering than they would have endured had the planned infrastructural changes been implemented. And the overall amount of animal suffering produced by those factory farms is greater than it would have been had those changes been adopted (in which case a larger number of animals would have endured less severe suffering).

Backfire 3 illustrates a further way in which a person's vegetarianism might do harm. But again we must ask: Is the prospect of one's vegetarianism leading to such harm merely a remote possibility that can be easily dismissed?

Of the consequentialist proponents of vegetarianism, Peter Singer most extensively engages with the general concern raised by *Backfire 3*. He doubts that the adoption of more humane animal-rearing practices would produce meaningful reductions in animal suffering. In part, this is based on his supposition that "[i]t is not practically possible to rear animals for food on a large scale without inflicting considerable suffering" (Singer 2009[1975]: 160). In addition, he cautions that raising animals for food at all, even if more humanely than in the status quo, encourages people to regard animals as mere objects. On this basis, he worries that if humane reforms were adopted, it would be nearly impossible over time for people to resist the tendency to cut costs by reverting to crueler methods (Singer 2009[1975]: 159–160, 229; Singer 1980: 332). Furthermore, he finds it implausible that reduced demand for meat brought about by vegetarians would make the industry less likely to adopt more humane animal-rearing methods. Quite the opposite, in his view. Pointing to unsuccessful efforts at meat industry reform in Britain, he concludes that meaningful reform will not arise unless the industry is forced to reckon with lower demand for its products (Singer 2009[1975]: 161–162).

So, it looks like Singer has a fairly straightforward basis for dismissing the worry about a person's vegetarianism doing harm in the way that occurs in *Backfire 3*. However, he does not cite the sort of evidence needed to establish his claims about vegetarianism proving a superior means of addressing animal suffering than efforts

focused on industry reform.¹⁴ Consider his claim that because the gains of reduced animal suffering that could be achieved via industry reform would be modest, our efforts would be better focused on abolishing the industry wholesale. Although it is plainly true that much greater gains would be achieved by abolishing rather than reforming the meat industry, it does not follow that our efforts should focus on the former goal over the latter goal. To draw that conclusion, we would need to consider projected trends concerning global meat consumption in the coming decades. And we would need to assess the related matter of how many current meat-eaters seem inclined to give up meat entirely versus being willing to pay higher prices for more humanely sourced meat. If we could reasonably expect a large segment of the population to continue eating sizable quantities of meat for years to come despite working to resist that trend, then efforts devoted to making farming practices more humane might hold greater promise in reducing net animal suffering after all.

Singer's concern that progress via industry reform would be easily reversible also turns on empirical matters that he does not sufficiently consider. For instance, it would be instructive to consider evidence on the frequency with which progressive reforms to farming practices (or to industry practices, more generally) adopted in the past have been overturned or weakened over time, and why that has happened when it has. Assessing his claim that meat-eating almost certainly would be accompanied by a tendency for humans to treat animals cruelly to save costs would require examining the prospects of a stable shift in public morals that might translate into ongoing support for more humane treatment of farmed animals.

The foregoing considerations would be beside the point if Singer were correct that vegetarians' efforts to reduce demand for meat are more likely to help than to hinder the cause of industry reform. But is he? The evidence he cites in support of this claim—the British case in which unsuccessful attempts at industry reform coincided with demand for meat remaining steady—hardly resolves the matter. All it shows is correlation, not causation. To help us further assess his claim, we might consider what has accounted for the adoption of industry reforms geared at more humane animal treatment where they have been adopted. In recent years in the U.S., a number of states have passed legislation aimed at making farming practices more humane, and some of the country's largest fast-food chains have made commitments to more humane supply chains in their animal products (Kaufman and Greenwood 2015; Norwood and Lusk 2011: 54–59; Strom 2012). What role, if any, has reduced demand for meat played in generating public support for such measures? It is not clear that it has played much of a role at all. As support for more humane practices has risen significantly, trends in meat consumption have remained more or less steady (Norwood and Lusk 2011: 58). Some have chalked up the ballot successes primarily to the savvy advocacy campaigns of animal welfare activist groups (Norwood and Lusk 2011: 54–59). And McDonald's commitment to more humane animal treatment, which is expected to translate into appreciable long-term

¹⁴ To be clear, I am not accusing Singer (or other consequentialist proponents of vegetarianism) of failing to engage with empirical evidence. These authors tend to do so extensively. Nevertheless, for their argument to succeed, their empirical case must be strengthened.

reductions in animal suffering, may be explained by the growing popularity of rival food chains, like Chipotle, that strive to use humanely sourced meat (Kaufman and Greenwood 2015). Indeed, the meat industry might not be incentivized to produce more expensive, humanely sourced meat unless it could count on robust demand for such products from meat-eaters.

Backfire 3 demonstrates that for the consequentialist argument for vegetarianism to be convincing, it will not do to merely point to the role that one's vegetarianism might play in reducing the number of animals reared in factory farms. In addition, an assessment of the expected consequences of different means of reducing animal suffering (including, but not limited to, industry reform aimed at more humane animal treatment) and their potential to be at odds with one another is needed. Carrying out such an assessment is no easy matter. In part, this is because it might not even be apparent to us when such harm occurs. Consider that if the scenario described in *Backfire 3* happened in the real world, it would likely seem to most of us that the efforts of vegetarians were successful since they reduced factory farming output. To see that those efforts, in fact, did harm, would require making a counterfactual judgment that relies on hard-to-attain information. The broader public might well never find out about nearly-adopted reforms and thus would have no way of knowing that greater reductions in animal suffering would have obtained if not for the reduced demand for meat caused by vegetarians. And if we cannot easily detect such harm occurring, then surveying similar historical cases in an attempt to assess the frequency with which past boycotts have inadvertently backfired by impeding other means of advancing a cause is unlikely to be informative. It might often go unnoticed when boycotts have been counterproductive in this way, and, indeed, failed boycotts may well be miscategorized as victories. However, until such issues are given due consideration, we face a further dimension of radical uncertainty concerning whether vegetarianism is justified based on its expected consequences.

5 The True Costs of Doing Good

A number of prescriptions for individual actions aimed at addressing large-scale harms have been, or might be, thought to follow from consequentialist considerations similar to those adduced in favor of vegetarianism. For instance, it has been argued that a person should purchase products that are both fair-trade and not made in sweatshops, over conventional alternatives, as long as doing so would not pose a great financial burden (Almassi 2011: 404). Such consequentialist reasoning has also been used to argue that one should refrain from greenhouse-gas emitting activities that are easily avoidable, such as driving a car when it's viable for one to instead cycle or use public transportation (Almassi 2012; Hiller 2011). Like the case of factory farming, these cases involve massive problems that demand institutional remedies. And in these cases too, given the scale of the harms at issue, it might seem initially doubtful that any one person's choices could make a meaningful difference. For example, it hardly seems that my decision to bike rather than drive to work could help prevent any future tsunami (Sinnott-Armstrong 2010), or that by

purchasing this pair of sneakers over another I could help improve working conditions in a distant country.

Consequentialist proponents of such actions acknowledge that we cannot determine with any precision the expected good of a person's acting in these ways. Still, they contend that the best available evidence indicates that a person might make *some* difference in addressing the large-scale harm at issue by acting in the prescribed way (Almassi 2011: 404 and 2012: 13–17; Hiller 2011: 357–358, 365–366 and 2014: 20–27, 33–34). And if one were to make a difference by acting these ways, the gains produced would be morally significant. I could crucially help reduce emissions by the exact margin needed to prevent a global temperature rise that would cause a greater incidence of catastrophic extreme weather events. Likewise, my refusal to purchase sweatshop-made apparel could prove pivotal to bringing about improvements to the miserable labor conditions currently endured by many garment industry workers. Although improbable, the potential good that such actions might produce is significant—and, some argue, plainly outweighs the relatively negligible personal sacrifices that performing such actions usually involves (Almassi 2011: 404 and 2012: 16, 18–19; Hiller 2011: 352, 356). Parallel reasoning seems to extend to other cases, telling in favor of individuals engaging in political-activist efforts to help combat serious injustices, e.g. by contacting one's representatives or circulating petitions. By giving up a few hours of one's free time each week, such actions could be crucial to reaching a critical mass of support needed to spur the adoption of important legislation. Thus, simple cost-benefit reasoning as that advanced in the case of vegetarianism seems to provide a straightforward explanation for why a person should act in a number of other specific ways too.¹⁵

However, the same sorts of problems I raised for the cost-benefit analysis used to support vegetarianism arise for the cost-benefit analyses used to support these other prescriptions. Consider how other well-intended efforts to address large-scale harms might backfire. Affluent consumers' refusal to buy sweatshop-made goods could cause sweatshops to close, and, in turn, thousands of workers to lose desperately needed jobs, while failing to produce any improvements to the objectionable labor practices that motivate such boycotts (Elgot 2013). Growing consumer support for fair-trade coffee might depress the already meager wages of the great majority of coffee-growers in developing nations who cannot afford to become fair-trade certified (Narlikar and Kim 2013). The rising popularity in almond milk spurred by opponents of dairy farming might have played a role in the aggressive expansion of almond production in California, exacerbating the state's already serious drought given the huge amounts of water that crop requires (Goodyear 2015). The seemingly eco-friendly choice to use reusable mugs over Styrofoam cups for one's hot beverages might be worse for the environment (Leibenluft 2008). By some accounts, the American government's 2010 initiative aimed at reducing the market for conflict minerals, which was adopted at the urging of social activists, has worsened the

¹⁵ The case for individuals withholding support for harmful conventional practices (e.g. by not purchasing sweatshop-made clothing) and for supporting positive alternatives (e.g. by purchasing sweat-free sneakers) both seem defensible by appeal to such reasoning.

already dire circumstances of many Congolese citizens (Aronson 2011). In decreasing its reliance on coal as an energy source in the past decade, as called for by environmentalists, the U.S. has reduced its domestic carbon emissions but has also rapidly increased its coal exports as a result, which some worry may lead to a net rise in global carbon emissions (Grose 2013).

The backfire cases just described are not especially anomalous. Individual efforts to address large-scale harms can and do backfire with non-negligible frequency. Thus, these cases are unlike *Crosswalk*, where I have sufficient background knowledge about like cases to conclude that the risk of backfiring is low enough to set aside. Still, it might seem that the backfiring outcomes in the surveyed cases were entirely unpredictable. And surely, our practical decision-making about whether to perform particular actions aimed at addressing large-scale harms hardly need take into account the merely theoretical possibility of those actions backfiring in utterly unpredictable ways.

However, the concern here is not about possible harms that are too remote to bear on our practical decision-making. The sorts of backfire scenarios I have described are *not* entirely unpredictable. Although I have underscored the considerable complexity of the causes of large-scale harms and the social contexts in which they occur, these complexities are not wholly beyond our ken. On the contrary, we know enough about them to know some of what we *don't* know, such as whether the sorts of backfire scenarios I have described are more or less unlikely to result from the actions in question than the more optimistic scenarios those actions are intended to produce. We know the sorts of information that would tell us what we need to know. And it is often clear enough what the empirical investigation needed to reveal such information would look like. The problem is just that, in many cases, we do not currently possess the requisite information. Thus, we cannot justifiably draw the sorts of conclusions that consequentialists think we can, about certain actions being warranted.

As before, I am not suggesting that such actions are more likely to do harm than good. And I acknowledge that in addition to actual cases of backfiring, there are actual cases of such actions doing good. My aim in surveying actual backfire cases is to buttress my premise that numerous individual actions aimed at addressing large-scale harms have the epistemic profile of *Virus* rather than *Crosswalk*. In many such cases, a person can conclude based on available evidence that by performing such an action, she might do great good or might do great harm, but it is at best unclear which outcome is more likely.

I cannot definitively show how far-reaching these problems are. Instead, I will describe one further complicating factor. Consider the following cases.

Backfire 4. Suppose that by replacing his energy-inefficient SUV with a compact hybrid vehicle, Adam will considerably lower the emissions that he produces. But because his monthly car payments are much higher than before, he can no longer afford to purchase the more humanely produced but pricier free-range eggs he used to buy, and so he switches to caged eggs. It turns out that Adam's decision to purchase the more fuel-efficient car has no positive environmental effects. Moreover, had he not bought the hybrid, and, conse-

quently, kept purchasing only free-range eggs, then demand for conventionally produced eggs would have fallen by just enough to result in significantly fewer egg-laying hens enduring horrendous, cramped conditions in the future. Alas, because he opts for the hybrid, this reduction in animal suffering does not come about.

Backfire 5. Previously, Adam and like-minded others had been gung-ho about doing everything they could to address the evils of factory farming—organizing PETA rallies, circulating petitions, staging protests of local vendors, and, of course, eating only vegan fare. But, worried about being too narrowly focused on just this one cause, they decide to broaden their horizons and seek to do good on many different fronts: buying fair-trade products and local produce; biking instead of driving; meticulously sorting recyclables and composting; organizing rallies for immigrants' rights and against workplace gender discrimination; and so on. Unfortunately, it turns out that had these individuals continued to devote themselves exclusively to fighting factory farming, they would have helped reduce animal suffering. By opting to address a variety of different causes instead, their efforts end up being spread too thinly, and they fail to make a difference on any of those fronts.

The problem these cases highlight concerns the opportunity cost associated with many efforts to do good. There are numerous large-scale harms, and often there are numerous, different means of addressing each of them. Because each person has finite moral resources—there are limits to the amount of time, money, and energy one can devote to doing good—working to address large-scale harms in some ways might constrain one's ability to do good in others. This consideration reveals a further way that attempts to address large-scale harms might backfire: as occurs in *Backfire 4* and *5*, a person's well-intended efforts might make no positive difference while preventing one from acting in other ways that would have done good. As we have seen above, those who argue for such actions by appeal to consequentialist reasoning characterize the costs of those actions in terms of the modest personal sacrifices they tend to imply. But, that construal of the costs involved is too narrow insofar as it ignores the matter of opportunity costs, a further factor that is relevant.

To assess the risk of one's efforts backfiring in this way, we would need to make comparative judgments about the expected utility associated with each of the many different competing options one might pursue to address large-scale harms. We would need a basis for determining whether one could expect to do greater good by spending her money on a new hybrid vehicle or free-range eggs (or, fair-trade coffee, solar paneling, campaign contributions, etc.). That, in turn, would require considering evidence concerning the odds of doing good as well as the amount of good one could expect to do by acting in the various ways. It seems reasonable to expect that resolving these matters would require in-depth investigation of complex empirical issues: issues similar to those that arise in working out the approximate expected payoff of vegetarianism. One such complexity involves the dynamic nature of some of the relevant variables. For instance, which causes my efforts might do the greatest good in advancing will partly depend on how much support there currently is for different candidate causes. And that issue, in turn, may be prone to sudden

fluctuation, based on hard-to-predict trends and events—ranging from the more mundane (e.g. which book Oprah happens to be reading this week or which TikTok video has recently gone viral) to the more serious (e.g. shifts in political leadership or disruptive changes to the global economy). Once again, the case for specific prescriptions grounded in consequentialist reasoning (say, for buying a hybrid, which might entail foregoing free-range eggs) rests on an empirical case that has not yet been adequately made. At least for now, sufficient justification for drawing the requisite conclusions about the relative expected payoffs of the various competing options for how one might address large-scale harms is lacking.

Another complication concerns how we are to determine which individual actions aimed at addressing large-scale harms deserve consideration in the first place. The actions that have been under discussion so far bear a relatively straightforward connection to the respective large-scale harms they are intended to address. However, as the arguments of the previous two sections suggest, the causal connections between a person's actions and that which might be needed to address large-scale harms are often not so straightforward. For instance, recall why things go wrong in *Backfire 2*: in that case, one's vegetarianism worsens polarization, and that in turn leads to a rise in meat production. And as we have seen, that imaginary case has real-life analogues; it cannot be dismissed as fantastical. For all we currently know, a person could better help to reduce animal suffering by spending more time socializing with people who have political views starkly opposed to her own rather than on efforts more straightforwardly linked to that cause, like campaigning for PETA (Jilani 2018). Despite the fact that the former activity bears a less obvious connection to the cause of reducing animal suffering, it may be that one could do more to further that cause by helping reduce polarization than by embracing what might seem the more obvious choice. Indeed, perhaps addressing many large-scale harms, ranging from animal abuse to global warming, will turn on the development of stronger bonds of solidarity between those who, at present, do not see eye to eye on those issues.¹⁶ And if this is so, then we might be unable to make meaningful progress without first reducing social division. Whether individual actions that are more directly connected to large-scale harms are more likely to do good than other options is an open question: one that should be resolved by engaging with relevant empirical evidence. Without doing so, it seems arbitrary to privilege actions that initially seem to bear a more straightforward connection to the harms at issue over less obvious alternatives.

The foregoing reflections point to the need for a *systematic* approach to the matter of how individuals should address large-scale harms. That is, instead of considering the case for any one specific action in isolation, we should consider how, across the numerous large-scale harms that confront us, a person might best direct her finite moral resources. To be sure, consequentialist proponents of such actions do recognize the importance of finding a reasoned basis for working out how individuals ought to adjudicate competing moral demands. However, these authors show less appreciation for the fact that this matter undergirds a further dimension of radical

¹⁶ Consider how some efforts to promote greater public awareness about climate change appear to have backfired due to deepening social division on that issue (Nordhaus and Shellenberger 2014).

uncertainty, and thus threatens to undermine consequentialist arguments for a variety of specific individual actions aimed at addressing large-scale harms.

Let me clarify the scope of my argument. It is not my contention that every individual action aimed at addressing a large-scale harm is prone to backfiring in each of the ways discussed above. For instance, not all ways in which a person might seek to address large-scale harms carry an opportunity cost—indeed, undertaking some such actions might expand, rather than constrain, one's capacity to do good. Nor do I deny that there might be some specific individual actions meant to address large-scale harms that escape the worries I have raised: actions whose performance is justified based on what we know at present. My aim has been to emphasize the great complexity involved in assessing the expected consequences of such actions—complexity that implies that we would have to know much more than we presently do to draw firm conclusions about how one should act in many, though not necessarily all, cases of this kind.

6 Conclusion

I have argued that consequentialist reasoning that initially might seem to provide clear practical guidance concerning how individuals should go about addressing many large-scale harms, upon reflection, fails to do so. Consequentialist arguments for particular such actions, such as embracing vegetarianism, often depend on an overly simplified construal of the costs and benefits associated with those actions. Proponents of those arguments fail to take seriously enough the prospect of those actions doing harm—regarding that as a merely remote possibility that can be put to the side. The rationale for doing so mistakes the *radical* uncertainty that characterizes our current epistemic situation with respect to addressing many large-scale harms for *ordinary* uncertainty of the sort illustrated by *Crosswalk*. The burden of showing that specific individual actions in these cases are justified in light of the various and sundry ways they might backfire is substantial, and it is considerably more weighty than consequentialist proponents of such prescriptions tend to acknowledge. I have argued that this burden has not yet been met in numerous cases of this sort: for individual actions that have been explicitly defended, such as vegetarianism, as well as for many others that might be defended by appeal to parallel consequentialist reasoning.¹⁷

If my argument is sound, then we often do not know how we should go about helping to address many large-scale harms. To some, that conclusion will likely seem incredible. But that is not surprising. It is a deeply unsettling fact that we are inextricably connected to, and often benefit from, immense suffering avoidably endured by others. Acting in concrete ways that we believe may help combat this suffering—by, say, refusing to eat meat or boycotting sweatshop-made products—can

¹⁷ My arguments have no bearing on defenses of such actions that are grounded in non-consequentialist considerations, such as the symbolic value of protesting unjust practices. For discussion of such views, see Hill (1979); Kutz (2000); Lichtenberg (2014).

feel empowering and even serve to alleviate some of the guilt associated with this troubling state of affairs. Be that as it may, given the present state of knowledge, to persist in the belief that we are meaningfully addressing such suffering by acting in those sorts of ways might be an instance of wishful thinking. Indeed, a potential danger of individuals engaging in such well-intentioned efforts is that they might feel they have thereby fulfilled their responsibilities and be less open to exploring more constructive alternatives. Thus, admitting that much more would have to be known to assess how ordinary individuals can meaningfully address many large-scale harms might just be an important initial step in tackling these pressing issues.

Acknowledgements For helpful feedback I am grateful to Torin Alter, Michael Horton, two anonymous referees for this journal, and audiences at the Australasian Association of Philosophy's annual conference in Dunedin, the Alabama Philosophical Society's annual conference in Pensacola, the University of Alabama, the University of Southern Mississippi, Arizona State University, and the University of Delhi. This work was supported by a grant from the University of Alabama's Research Grants Committee.

References

- Almassi, Ben. 2011. The consequences of individual consumption: A defence of threshold arguments for vegetarianism and consumer ethics. *Journal of Applied Philosophy* 28(4): 396–411.
- Almassi, Ben. 2012. Climate change and the ethics of individual emissions: A response to Sinnott-Armstrong. *Perspectives: International Postgraduate Journal of Philosophy* 4(1): 4–21.
- Aronson, Dan. How congress devastated congo. *The New York Times*, August 7, 2011.
- Bollyky, Thomas J. The global grip of cigarette smoking. *The Los Angeles Times*, Jan 21, 2014.
- Champene, Aaron, and Don Merrell. 2008. The causal impotency objection to vegetarianism. *Southwest Philosophy Review* 24(1): 53–60.
- Chartier, Gary. 2006. On the threshold argument against consumer meat purchases. *Journal of Social Philosophy* 37(2): 233–249.
- Cole, Matthew, and Karen Morgan. 2011. Vegaphobia: Derogatory discourses of veganism and the reproduction of speciesism in UK national newspapers. *The British Journal of Sociology* 62(1): 134–153.
- Doering, Christopher. U.S. to buy \$170M in meat, fish to help farms hit by drought. *USA Today*, August 13, 2012.
- Elgot, Jessica. Bangladesh factory collapse: Should you boycott primark over workers' rights? *The Huffington Post*, June 5, 2013.
- Feldman, Fred. 2006. Actual utility, the objection from impracticality, and the move to expected utility. *Philosophical Studies* 129(1): 49–79.
- Foran, Clare. America's political divide intensified during trump's first year as president. *The Atlantic*, October 5, 2017.
- Friedman, Monroe. 1999. *Consumer boycotts: Effecting change through the marketplace and the media*. New York: Routledge.
- Glickman, Lawrence B. 2009. *Buying power: A history of consumer activism in America*. Chicago: Chicago University Press.
- Glickman, Lawrence B. Chick-fil-a day a reminder that boycotts often backfire. *Bloomberg News*, August 3, 2012.
- Goodyear, Dana. What milk should i drink? *The New Yorker*, October 23, 2015.
- Grose, Thomas K. As U.S. cleans its energy mix, it ships coal problems abroad. *National Geographic*, March 17, 2013.
- Hill, Thomas E. 1979. Symbolic protest and calculated silence. *Philosophy & Public Affairs* 9(1): 83–102.
- Hiller, Avram. 2011. Climate change and individual responsibility. *The Monist* 94(3): 349–368.
- Hiller, Avram. 2014. A "famine, affluence, and morality" for climate change? *Public Affairs Quarterly* 28(1): 19–39.
- Hudson, Hud. 1993. Collective responsibility and moral vegetarianism. *Journal of Social Philosophy* 24(2): 89–104.

- Jilani, Zaid. A New Study Shows how american polarization is driven by a team sports mentality, not by disagreement on issues. *The Intercept*, April 3, 2018.
- Kagan, Shelly. 2011. Do i make a difference? *Philosophy & Public Affairs* 39(2): 105–141.
- Kaufman, Alexander C., and Arin Greenwood. McDonald's plans to switch to cage-free eggs. *Huffington Post*, September 9, 2015.
- Kutz, Christopher. 2000. *Complicity: Ethics and law for a collective age*. Cambridge: Cambridge University Press.
- Leibenluft, Jacob. Grande americano, extra green: What's the best way to drink your morning joe? *Slate*, September 16, 2008.
- Lenman, James. 2000. Consequentialism and cluelessness. *Philosophy & Public Affairs* 29(4): 342–370.
- Lichtenberg, Judith. 2014. *Distant strangers: Ethics, psychology, and global poverty*. Cambridge: Cambridge University Press.
- Matheny, Gaverick. 2002. Expected utility, contributory causation, and vegetarianism. *Journal of Applied Philosophy* 19(3): 293–297.
- Narlikar, Amrita, and Dan Kim. Unfair trade: The fair-trade movement does more harm than good. *Foreign Affairs*, April 4, 2013.
- New York Times. The global tobacco threat. *The New York Times*, Feb. 9, 2008.
- Nobis, Nathan. 2002. Vegetarianism and virtue: Does consequentialism demand too little? *Social Theory and Practice* 28(1): 135–156.
- Norcross, Alastair. 2004. Puppies, pigs, and people: Eating meat and marginal cases. *Philosophical Perspectives* 18: 229–245.
- Norcross, Alastair. 2008. Causal impotence and eating meat. *Southwest Philosophy Review* 24(2): 5–10.
- Norcross, Alastair. 2012. Puppies, pigs, and potency: A response to galvin and harris. *Ethics, Policy & Environment* 15(3): 384–388.
- Nordhaus, Ted, and Michael Shellenberger. Global warming scare tactic. *The New York Times*, April 8, 2014.
- Norwood, F.Bailey, and Jayson L. Lusk. 2011. *Compassion, by the pound: The economics of farm animal welfare*. Oxford: Oxford University Press.
- Rachels, Stuart. 2011. Vegetarianism. In *The oxford handbook of animal ethics*, ed. Tom L. Beauchamp and R.G. Frey, 877–905. Oxford: Oxford University Press.
- Regan, Tom. 1980. Utilitarianism, vegetarianism, and animal rights. *Philosophy & Public Affairs* 9(4): 305–324.
- Reynolds, George. Why do people hate vegans? *The Guardian*, Oct. 25, 2019.
- Shafer-Landau, Russ. 1994. Vegetarianism, causation and ethical theory. *Public Affairs Quarterly* 8(1): 85–100.
- Singer, Peter. 1980. Utilitarianism and vegetarianism. *Philosophy & Public Affairs* 9(4): 325–337.
- Singer, Peter. 2009[1975]. *Animal liberation*, revised edition. New York: HarperCollins.
- Sinnott-Armstrong, Walter. 2010. It's not my fault: Global warming and individual moral obligations. In *Climate ethics: Essential readings*, ed. Stephen M. Gardiner, Simon Caney, Dale Jamieson, and Henry Shue, 332–346. New York: Oxford University Press.
- Strom, Stephanie. McDonald's vows to help end use of sow crates. *The New York Times*, Feb. 13, 2012.
- Telegraph Reporters. 'We'd rather go back to the streets than eat vegan food' homeless tell famous Italian chef. *The Telegraph*, April 28, 2016.
- Tomson, Bill. U.S. Makes special purchase of \$40 million of chicken products. *The Wall Street Journal*, August 15, 2011.
- Yaffe-Bellany, David. Arby's has an answer to plant-based meat: A meat-based carrot. *The New York Times*, July 16, 2019.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.