

Intensive Animal Agriculture and Human Health

Jonathan Anomaly
Routledge Handbook of Animal Ethics
Edited by Bob Fischer
Forthcoming, 2019

This chapter provides an overview of some ethical issues that arise when people use intensive agricultural techniques – often called “factory farming”¹ – to raise livestock for food. I will ignore important issues like animal suffering and environmental concerns that other authors in this book discuss. Instead, I will focus on problems factory farming raises for human health, with special emphasis on antibiotic-resistant bacteria that can arise in animals and infect people. I conclude that there are legitimate disagreements about what we should do about factory farming, but I argue that everyone should agree on two things: first, we should defer to the best available science rather than viewing these issues through the prism of political ideology; second, we should think carefully about the moral tradeoffs of different policy proposals rather than pontificating from an armchair.

1. Introduction

The division between industrial farms and small farms that use traditional techniques is not a clear one. Most animal farms fall somewhere between large and small, and most farmers use varying degrees of confinement to house their animals. More importantly, the distinction between industrial farming and traditional farming does not track what is morally wrong and right. It is commonly assumed that if there are moral problems with industrial farming, it must be due to the size and scale of the farms. But this view cannot be right, since many companies (like Whole Foods and Chipotle in the United States) source their products from large farms with morally praiseworthy practices (Robbins *et al* 2016).

In fact, the high standards and low prices offered by purveyors of milk and meat and eggs are often made possible by the large scale of the operations from which they buy their food. Moreover, contrary to popular thought, small farms can be worse for the environment than large farms because of the inefficiency of transporting small amounts of food and equipment back and forth (Desrochers and Shimizu 2012), and because sometimes even *small* farms employ techniques like extreme confinement which requires them to use antibiotics to prevent infections in their livestock. Finally, small farms in which chickens are left free to roam outside are more likely to expose their animals to pathogens they pickup from wild birds (especially avian flu), which can harm the farm animals themselves, and the people who eat them (Greger 2007).

In other words, we should not think of the moral differences between practices as corresponding in any way to the size of the farm or the novelty of technology used in farming.

¹ I use “factory farmed” to imply the extreme confinement of animals, and the routine use of antibiotics. Factory farming is a vague term, but it typically includes these two elements.

Technology is a tool: it can be used to improve the lives of people and animals, or make them worse. What makes intensive animal agriculture such an interesting moral issue is that there are no easy solutions to the problems it poses. Replacing animal protein with plant protein, or creating artificial meat from stem cells, are not yet viable alternatives for many people around the world. Until everyone has access to cheap protein that doesn't come from animals, we need to consider how to improve current farming techniques that threaten animal welfare and human health.

2. The Science of Antibiotic Resistance

Antibiotics have probably saved more lives in the twentieth century than any other medicine. It's possible that vaccines against viruses that cause debilitating diseases like polio and smallpox have *prevented* more diseases than antibiotics have *cured*, but either way it is worth appreciating the enormous amount of suffering and death antibiotics have helped to prevent. In addition to their ability to cure bacterial infections, antibiotics also make surgery far safer than it would otherwise be. Before the discovery and mass production of antibiotics, a skin laceration or an infected tooth could be lethal because of the opportunity this gave for pathogenic (disease-causing) bacteria to enter the bloodstream and cause organ damage or system failure. The same thing is true for farm animals, though antibiotics are mainly used on factory farms to promote the growth of animals and prevent infections that animals are prone to contract in crowded and stressful conditions (Anomaly 2015).

Most bacteria are not harmful, and in fact some are helpful to their hosts. Of the approximately 40 trillion bacteria living on and in our body at any given time (Sender 2016), many are *commensal*, meaning that they use us as a source of food and shelter but do not harm us, and some are *mutualistic*, meaning that they provide us with positive benefits in exchange for food and shelter. These benefits can include priming our immune system to distinguish friend from foe, helping us digest food, modulating our hunger, and even altering our mood (Velasquez-Manoff 2013, McAuliffe 2017). Of course, there is no conscious exchange going on. Mutualistic relationships often emerge between organisms when they interact over long periods of time (Trivers 1971).

Parasitic relationships occur when bacteria thrive at the expense of their hosts. For example, the bacteria that cause tuberculosis can impair our ability to fight off other infections, and can kill us when we are especially vulnerable. Some bacteria provide benefits in small concentrations (usually by crowding out pathogenic bacteria) but costs in large numbers, especially when they enter our bloodstream or invade parts of our body where they are not normally found (Blaser 2014). In fact, *mycobacterium tuberculosis* may be an example: some evidence suggests that the complete *absence* of bacteria that cause tuberculosis increases our risk of developing multiple sclerosis, even if the same bacteria can kill us when we have a secondary infection or our immune system is weak (Velasquez-Manoff 2013). This is part of a more general pattern in which eliminating microbes we have co-evolved with for millions of years can elevate the risk of developing autoimmune diseases, which occur when our immune system searches for a target that it cannot find and attacks our own cells instead. Our relationship with bacteria is complicated.

Antibiotics are a crucial weapon in fighting off harmful bacterial infections. But the more we use antibiotics, the less effective they tend to become. This is because the creatures that bacteria parasitize have evolved weapons to kill or disable them, and bacteria have evolved

mechanisms to resist the effects of these weapons. For example, some antibiotics (produced naturally by living creatures or artificially in the lab) penetrate bacterial cell walls, or disrupt the ability of bacteria to replicate. Bacteria often respond by creating thicker cell walls or pumps that eject the antibiotics they absorb. This evolutionary arms race has been going on for billions of years, and all we can do is alter its course by devising new antibiotics, new vaccines, and deciding how to use existing antibiotics responsibly.

According to the most comprehensive analysis yet produced (O'Neill 2015), a consensus of researchers believes that there is a causal link between the use of antibiotics in agriculture, and the rise of antibiotic resistance in human populations. It might seem odd that *all* scientists don't agree that agricultural use causes resistance in people. But scientists are cautious in separating correlation and causation. What we can show is that in places where the use of antibiotics in agriculture is high, resistance is also high, and when countries cut back on their use of antibiotics in agriculture (for example, European Union countries over the last two decades), resistance tends to decline (Spellberg 2016). But it's also true that in places where antibiotic use in animals is high, it tends to be high in people, which means resistance is generally high in those countries in both people and animals. Since resistant strains of bacteria that arise in people can be transferred to animals, it is difficult to determine with certainty what the causal direction of a particular pattern of resistance is: it can spread from people to animals or animals to people.

Of course, people *are* animals, and we share a bacterial environment with non-human animals. Not all bacteria that are dangerous for other animals cause problems in people, or vice-versa, so some scientists think we should worry more about using antibiotics in agriculture that are medically important for people (e.g. penicillin), rather than antibiotics in general. But since bacteria swap genes with one another, and can acquire bits of DNA from bacteria of any species, when we encourage antibiotic resistance in a kind of bacterium that doesn't harm human health, the genes that confer resistance may very well be transferred to other species of bacteria that threaten human health (Marshall and Levy 2011). Thus, the more we use antibiotics in animals, even antibiotics that fight bacteria that aren't harmful to human health, the more we encourage the spread of antibiotic resistance in our general microbial environment. This is an inevitable consequence of evolution by natural selection.

It may be worth mentioning in passing, at least, how devastating the diminished efficacy of antibiotics is for human health. According to a recent analysis by a team of social and natural scientists, at least 700,000 people die each year of antibiotic resistant infections worldwide, and many people whose infections are eventually cured have to spend more money on second and third generation antibiotics, which are more expensive, because bacteria are becoming increasingly resistant to older antibiotics (O'Neill 2016). Since in most countries the costs of treatment are at least partly borne by taxpayers or members of an insurance pool, they are not fully felt by the individual, and are therefore largely ignored. But these costs are growing quickly, and we have good reasons to believe that antibiotics in agriculture plays some role in this, even if it is not (yet) a bigger role than the direct use of antibiotics in people.

3. Ethics and Antibiotic Resistance

a. consent

One reason the spread of antibiotic resistance from farm animals to people is a *moral* problem is that it constitutes a harm to which nobody has an opportunity to consent. Adding more resistant

bacteria to our shared microbial environment can be thought of as a harmful byproduct, or negative externality, of antibiotic use. Resistance rises whether the antibiotics are used to treat an infection, promote growth in farm animals, or are used for no good reason. Bacteria don't care why we use antibiotics. They just tend to respond to antibiotics in their environment by increasing their resistance (more accurately, bacteria without resistance genes die off and make way for their resistant cousins, which harm people because they are more difficult and costly to treat).

In this way, antibiotic resistance is a lot like air and water pollution: even those who don't contribute to pollution bear the costs. Most people who produce pollution don't intend to harm anyone, most victims aren't given a chance to consent to the dangers of pollution, and when individuals or companies are permitted to pollute without compensating victims, we tend to get more pollution than is socially desirable. In the language of economists, we get an "inefficient" amount of pollution when producers and consumers fail to internalize the costs of their transactions.

Collective consequences aside, focus on the following question: is it ever morally permissible to inflict harms on people without their consent? An obvious case in which this seems justifiable is a parent overriding their children's desires in the short run in order to provide long run benefits. When I tell my daughter, Loki, that she'll only get cake after she does her school work, I frustrate her desire for instant gratification, and in this sense I harm her. But in cases like this I'm presumably justified in my *belief* that this will produce a better life for her, and justified in my *action* of establishing penalties and incentives for getting her to do the right thing. The hope is that Loki will thank me later. I have good reasons to believe she will because she'll eventually recognize that children have poor impulse control and benefit from developing healthy habits. Even if she thanks me later, though, recognizing why I restricted her liberty doesn't constitute retroactive consent. Nor is it important that, when the child becomes an adult, she explicitly thinks about why her parents' rules benefited her. This suggests that while consent is morally important, it is not the only thing that matters in justifying our actions to other people, even when those actions inflict harm on them.

Of course, farmers adding antibiotics to animal feed are different than parents withholding cake from kids. Antibiotic resistance can be thought of as genetic pollution dumped in the broader microbial environment, which poses a small probability of endangering a large number of people. Unlike parents restricting their children's liberty to promote their welfare, many farmers do not *know* or *care* as much about the people whose welfare they indirectly affect. So even if there are cases in which one person is justified in overriding another's liberty to promote their welfare, without seeking their consent, it is not clear that the use of antibiotics in agriculture is such a case.

b) responsibility

Our common moral practices are infused with the assumption that responsibility matters. We praise or punish someone because they *deserve* rewards or penalties, and they deserve their treatment to the extent that they are *responsible* for their behavior. Of course, sometimes people who seem to deserve punishment actually do not, and there are many kinds of exculpatory facts that might lead us to excuse someone for acting in a way they weren't fully responsible for. So, for example, many people think we should punish thieves. But if we find out that a thief was given an ultimatum by a neighbor to either steal a pepperoni pizza or take a bullet to the head, we might excuse the thief because he was under duress.

There are disagreements about the justifications for punishment and praise, but on most accounts responsibility is crucial. Who is responsible for antibiotic resistance when it arises as a byproduct of voluntary choices by farmers and consumers? Responsibility often involves causal judgments about who produces an outcome that we think of as bad or good. So, for example, we might say farmers who use practices that produce genetic pollution (antibiotic resistance), along with consumers of meat from animals fed antibiotics, are *responsible* for antibiotic resistance. In one sense this is obviously true: both producers and consumers collectively *cause* the outcome we are concerned with, and in most cases they have the opportunity to behave otherwise.

Notice, however, that any one act of consuming factory farmed animals, or adding antibiotics to animal feed, does not necessarily cause antibiotic resistant bacteria to emerge or spread. All we get is a small probability that increases as the quantity of antibiotics involved in these practices grows. Thus, it is hard to say that any one action, or small set of actions, makes a difference (Nefsky, 2018). Still, we might say that each person who produces or consumes factory farmed meat bears some (perhaps probabilistic) responsibility for producing collectively bad consequences. This might not be a decisive reason to change our practices, but it may be a reason to think about how we might change the incentives that predictably lead to the negative consequences of factory farming.

Notice that there is nothing especially unique about antibiotic resistance. Many public health threats, including various kinds of air and water pollution, have a similar structure. Often chemicals that can threaten human health are benign below a certain threshold. For example, adding a teaspoon of mercury into the municipal water supply or a gallon of arsenic into the ocean will probably not produce any negative health consequences for people. But above some threshold, we may be contributing to a process that leads to adverse health consequences like an elevated risk of cancer in adults and neurological disorders in children. Here again, the typical response is not to punish people who are responsible for making a trivial contribution to a social problem. Instead, it is to establish regulations on the amount of pollution that can be produced in a particular industry. Of course, if this is true, it may shift responsibility onto all citizens to lobby governments to establish socially beneficial regulations (Sinnott-Armstrong, 2005).

The problem is that most citizens are predictably ignorant about many processes that constitute collective action problems – situations in which what is individually rational is collectively harmful. It may be rational for *each of us* to pollute, even if *all of us* would be better off if each polluted less. Economists use the phrase “rational ignorance” to cover cases in which people have little reason to gather information because the individual costs of doing so exceed the individual benefits (Downs 1957). The idea is that even if all of us would be better off becoming informed about a particular case of pollution, and voting for candidates that pledge to cost-effectively regulate it, each of us has an incentive not to be informed. This is because each of us has little power to do much about the problem even if we understood its precise causes. After all, consuming information is costly in the opportunity sense: we have to give up time and energy and set aside our ideological commitments to confront uncomfortable facts.

The upshot is that responsibility for *causing* the negative effects of intensive animal farming is diffuse, and so is responsibility for *addressing* it. Effectively solving the problem requires many of us to act against our self-interest. To the extent that each of us is responsible for producing the problem, perhaps each of us has a moral reason to gather information about the source of the problem and the set of feasible solutions, and then act on that information at the ballot box, in the market place, and in our ordinary conversations with people. This may seem like a trivial moral obligation. But becoming a responsible voter and consumer, and a

sophisticated interlocutor in political conversations, is demanding. As I'll argue in the next section, it requires us to become aware of our biases, cultivate skepticism, and face up to facts that collide with our prior beliefs.

4. The Ethics of Belief

In the previous section, I mentioned the concept of rational ignorance from economics, and deployed it to try to explain why most consumers and voters don't understand the details of many of the collective action problems they face, including antibiotic resistance. Since each of us barely contributes toward the problem, we face a question. Why gather any information at all if, once I become informed, I am unlikely to be able to make much difference? Unless I'm a powerful politician, wealthy lobbyist, or influential in some other way, it seems difficult to justify the belief that my actions will matter much.

Another way to think about the issue is that there are three separate barriers to solving a collective action problem like antibiotic resistance. First, enough people have to gather information, even when there are few if any benefits for each of them to do so. Second, each must spend time sifting through the information in order to determine whether the relevant claims are true, whether they're morally important, and which policies or consumption habits are most likely to advance the cause of minimizing antibiotic resistance. Finally, each must *act* on the information by, for example, paying a higher price at the meat counter, convincing people to vote for candidates likely to pass reasonable regulations on antibiotics, etc.

Since antibiotic resistance is only one of many collective action problems, it is understandable, though not necessarily morally excusable, why many people would be ignorant about issues like the role they play in producing the problem. But the idea of rational ignorance understates the magnitude of the problem. As Michael Huemer (2015) argues, many people are *irrational*, especially in the political realm.

To understand the difference between rational ignorance and irrationality, we need to make a distinction: **epistemic rationality** aims at justification, while **instrumental rationality** aims at utility. When we seek to match our minds to the world by carefully considering evidence, we are behaving in an epistemically rational way. This sense of the term is captured when we say that a religious fanatic who rejects scientific explanations in favor of magical explanations is "irrational." Epistemically rational people proportion their beliefs to the evidence.

But in economics, and in everyday life, we also use the term "rational" to refer to someone who efficiently satisfies their desires. To take an example, we sometimes say of a brutal dictator that he is behaving in a way that is *wrong* but *rational*, if his goal is to stay in power. That is, his actions are morally wrong, but effective at achieving his objectives.

Putting these pieces together, as long as we have preferences about how we wish the world were, it can be instrumentally rational to render ourselves epistemically irrational. This is sometimes called rational irrationality, and it is pervasive in politics (Caplan 2001). Beliefs can function like placebos that provide us with meaning (e.g. religion), bond us together in a common cause (e.g. politics), or alter our confidence in ways that affect our performance (e.g. sports).

To take a simple example, since religious and political beliefs often function as tools of social bonding between communities, it can be painful, even devastating, to deviate from the beliefs your classmates, friends, or family members share. And it can be inspiring to believe that

whatever hardships we undergo in pursuit of our goals, there is a greater purpose that redeems our effort and suffering. Becoming an atheist, or declaring yourself a political heretic, can mean losing friends, missing out on social events, and even depriving yourself of a source of transcendent meaning. It is obvious, then, why people not only refrain from forming beliefs when they don't see a point in exerting effort (rational ignorance), but form strong beliefs they cannot possibly justify (rational irrationality). In fact, in political activism and religious practice, it is often a virtue to believe *in spite of* the evidence, rather than because of it. True believers are implored to have faith, not use reason.

Why are people especially prone to epistemic irrationality in politics? In addition to political beliefs being tools of social bonding for rival tribes, misguided political beliefs are rarely punished. In markets, consumers get what they pay for, so they are a little more careful to sift through information about the products they buy. When faced with alternative brands of car, consumers have to live with the consequences of what they choose. But in democratic politics, citizens get what *other people* vote for, so they are less prone to correct their errors of judgment. Voting one way or another doesn't *cause* an outcome to change, or a collective action problem to be solved. The fact that we rarely pay any substantial cost for forming irrational political beliefs may explain why otherwise smart people become good rationalizers but poor reasoners when thinking about politics (Haidt 2012).

If political ignorance and irrationality are predictably widespread features of life in a large democracy, it may be that our fundamental moral obligation as citizens is to resist the temptations of believing what we would like to be true, and deferring to experts on relevant matters. In other words, our main *moral* obligation may be to become *epistemically* rational, rather than succumbing to the irrationality that characterizes political discussions.

Committing ourselves to form beliefs on the basis of evidence cuts against the tendency of partisans on all sides. For example, it is common (in modern American politics) for self-described progressives to criticize religious conservatives for rejecting evolutionary explanations about the origin of human beings, or for dismissing the idea that human activity is partly responsible for recent changes in the earth's temperature. But the same people who profess an allegiance to science often reject it when it conflicts with their own political ideology. To take an example, skepticism about the safety of vaccines and genetically modified (GM) food more often comes from progressives than conservatives, despite a consensus that vaccines and GM food are generally safe and socially beneficial (NAS 2016).

If these considerations are right, the most fundamental moral obligation consumers and voters arguably have is the duty to proportion their beliefs to the best available evidence, and to defer to expert opinion when appropriate. Unless we recognize the obligation to form justified beliefs about politically-relevant subjects, it is not clear how we can hope to solve collective action problems like antibiotic resistance that arises in livestock and spreads to people.

5. Ethics and Public Policy

I've tried to show that the problem generated by the widespread use of antibiotics in agriculture is one instance of a set of problems that shares a common structure. Like many other global collective action problems, responsibility for producing the problem is diffuse, individuals cannot solve it on their own, and individuals have little incentive to carefully study the nature of the problem or the available solutions. Given this common structure, self-interest pulls us in one

direction and morality *should* pull us in another. But we cannot recognize our moral obligations as consumers or voters until we first overcome the strong temptation to form politically-motivated rather than epistemically rational beliefs about the nature of the problem.

It is likely that our sense of fairness, or justice, evolved in part to solve collective action problems in small groups (Bowles and Gintis 2013). In other words, we likely evolved moral emotions like guilt and shame because small groups of people with moral emotions can solve collective action problems more effectively than those who need to constantly monitor one another and impose physical costs on people who violate mutually beneficial norms or rules. The problem is that global collective action problems whose consequences spill across generations – like antibiotic resistance, or human-induced global warming – often fail to activate moral emotions that evolved to achieve much more modest goals in small groups. So one obvious question is how to convince individual farmers, consumers, and citizens to care about an outcome they have so little control over, but which has potentially devastating effects.

Fortunately, there is some evidence that consumers are willing to pay a tax on products that produce negative side-effects, if the tax revenue raised will be used to minimize the problem (Kallbekken 2011). And there is evidence that most American consumers want antibiotic-free meat, and are willing to pay at least a little more for it, even if many don't fully understand the problem (Consumer Reports 2012). Since ordinary people can't possibly compute the total costs and benefits their consumption choices produce, they cannot be expected to know how much they *should* be willing to pay for products that produce collective harms.

One thought, then, is that the primary moral obligation policymakers have is to try to implement policies that price pollution – understood broadly in this context to include antibiotic resistance – in a way that reflects what people *would* be willing to pay if they knew the relevant information. This is what is known in economics as a Pigovian tax, after the British economist Arthur Pigou, who argued that we should aim at an “optimal” pollution rate. The idea is that some pollution makes productive processes possible. For example, even extremely efficient factories typically emit a little pollution, and preventing or cleaning up the last few particles of pollution becomes increasingly costly relative to the benefits it produces (think of a beach cleanup where it takes many hours to find the last particle of trash: the time and energy could be better spent on other things that produce individual or social value).

In theory, Pigovian taxes reflect the precise social cost of an activity, and in addition to discouraging inefficient levels of pollution, the revenue generated is used to compensate victims of pollution, or invested in finding new technologies that produce similar products with less pollution. Since the information required to calculate a Pigovian tax with precision is virtually impossible to gather, many argue for less precise “user fees” on antibiotics that serve a similar function (Battin *et al* 2008, Anomaly 2013).

Notice that a Pigovian tax, or a slightly less precise user fee, serves another moral function apart from producing the socially optimal amount of pollution in the form of antibiotic resistance. That is, it doesn't just focus on the collective consequences of our activities, which some philosophers think of as a morally controversial approach to ethics. It also reflects a view of fairness according to which people should pay costs for acting in harmful ways in proportion to the amount of harm their activity can be expected to create. This principle of proportionality goes back at least to Aristotle. The idea is that we should pay for an activity in proportion to the amount of social costs we create by engaging in it, especially when we benefit at the expense of others.

Another way of characterizing this idea is to call it a public harm principle. According to a common (though not universally shared) interpretation, John Stuart Mill argued that people should be free to act as they please provided they don't create morally significant, uncompensated harm to others (1859). What counts as a harm and what counts as a compensating benefit is up for debate. But Mill's framework is meant as a general principle for framing moral debates, not an algorithm that provides us with detailed answers for moral questions. Some extend Mill's harm principle to large number cases (Gaus 1999). On this view, if a set of actions causes an uncompensated harm, and if the harm is serious enough that it merits limiting one person's liberty, then everyone's liberty should be similarly limited.

Moral principles are easier to apply between a small number of people than across a population of millions. One reason for this is that as the number of people involved increases, conflicting judgments about what should be done tends to grow, and solutions are more difficult to implement. When someone steals a car or attacks someone at a pub, it is clear who should pay the costs, and juries don't have too hard a time determining appropriate penalties. But anytime we think about political solutions to large collective action problems, we should assume that voters, policymakers and bureaucrats who implement policies will have imperfect information and conflicting incentives. In order to win election campaigns, policymakers often have to tell people what they want to hear, and we have already discussed why what people want to hear is not always what scientists know to be true.

So quite apart from the problem of politicians having limited knowledge, they do not always have incentives to use their knowledge to promote broader social goals unless voters can be expected to reward them for it. Political feasibility may therefore be a relevant constraint to take seriously in moral theorizing about how we should approach real world problems like antibiotic resistance in agriculture. When thinking about what to do, we should avoid complacently assuming insuperable voter ignorance and political corruption, on one hand, and perfectly informed and altruistic political actors, on the other.

One of the most important trends in antibiotic resistance is the precipitous rise in antibiotic use in agriculture in the developing world, which is occurring at the same time that antibiotic use in many developed countries is beginning to decline (van Boeckel *et al* 2015). In developed countries, the decline is due in part to demand by consumers for antibiotic-free meat, and to policymakers banning the use of antibiotics as growth-promoters in the European Union. Nevertheless, as global trade continues to increase in the coming century – a trend that mostly benefits consumers – there is a real danger that many factory farms will migrate to nations with weak regulations on the use of antibiotics in agriculture. Some have therefore called for global agreements to restrict the use of antibiotics in agriculture in order to avoid a situation in which countries with the weakest regulations sell products that produce global harms to people in countries with more responsible production practices (Anomaly 2018).

Conclusion

The widespread use of antibiotics in agriculture has only occurred for about half a century, and the consequences have been anticipated since the beginning. Alexander Fleming, who developed penicillin as a drug to treat serious bacterial infections in people, famously warned that its profligate use would lead to drug resistance, which he described as a *moral* problem. He never

anticipated that we might squander the efficacy of antibiotics even faster by feeding them to our farm animals in order to obtain relatively small benefits from cheaper meat.

Banning antibiotics in agriculture would likely benefit people by slowing antibiotic resistance. But it would also prevent farmers from treating animals who get sick even when the farmers engage in otherwise responsible production practices. I leave it up to the reader to decide whether we should opt for a complete ban or instead support some combination of taxes and regulations on the use of antibiotics in agriculture. The tradeoffs are difficult, and the benefits and costs of different approaches will be disproportionately borne by different people.

References

- Anomaly, Jonathan. 2013. Collective Action and Individual Choice. *Journal of Medical Ethics*, 39, 752-756.
- Anomaly, Jonathan. 2015. What's Wrong with Factory Farming? *Public Health Ethics*, 8(3): 246-254.
- Anomaly, Jonathan. 2018. Antibiotics and Animal Agriculture: The Need for Global Collective Action. In Michael Selgelid, editor, *Ethics and Antimicrobial Resistance*. New York: Springer Publishing.
- Battin, Margaret et al. 2008. *The Patient as Victim and Vector: Ethics and Infectious Disease*. Oxford University Press.
- Blaser, Martin. 2014. *Missing Microbes: How the overuse of antibiotics is fueling our modern plagues*. New York: Picador-MacMillan.
- Bowles, Samuel and Herbert Gintis. 2013. *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton University Press.
- Caplan, Bryan. 2001. Rational Ignorance versus Rational Irrationality. *Kyklos* 54: 3-26.
- Consumer Reports. 2012. Meat on Drugs.
http://www.consumerreports.org/content/dam/cro/news_articles/health/CR%20Meat%20On%20Drugs%20Report%2006-12.pdf. Accessed July 20, 2017.
- Desrochers, Pierre and Hiroko Shimizu. 2012. *The Locavore's Dilemma: In praise of the 10,000 mile diet*. New York: Perseus Books.
- Downs, Anthony. 1957. An Economic Theory of Political Action in a Democracy. *The Journal of Political Economy*, 65(2): 135-150.
- Gaus, Gerald. 1999. *Social Philosophy*. London, UK: M.E. Sharpe Publishing.
- Greger, Michael. 2007. The Human/Animal Interface: Emergence and Resurgence of Zoonotic Infectious Diseases. *Critical Reviews in Microbiology*, 33: 243-299.
- Haidt, Jonathan. 2012. *The Righteous Mind*. New York: Vintage Books.
- Huemer, Michael. 2015. Why People are Irrational About Politics. In *Philosophy, Politics, and Economics: An Anthology*. Edited by Jonathan Anomaly et al, Oxford University Press.
- Kallbekken, Steffen et al. 2011. Do you not like Pigou, or do you not understand him? *Journal of Environmental Economics and Management*, 62(1): 53-64.
- Marshall, Bonnie and Stuart Levy. 2011. Food Animals and Antimicrobials: Impacts on Human Health. *Clinical Microbiology Reviews*, 24(4): 718-733.
- Lane, Nick. 2015. *The Vital Question: Energy, Evolution, and the Origin of Complex Life*. New York: WW Norton and Company.
- McAuliffe, Kathleen. 2017. *This is Your Brain on Parasites: How tiny creatures manipulate our behavior and shape society*. Boston: Mariner Books.
- Mill, John Stuart. 1859. *On Liberty*. <http://oll.libertyfund.org/people/john-stuart-mill>
- National Academy of Sciences. 2016. *Genetically Engineered Crops: Experiences and Prospects*. Washington, DC: The National Academies Press.
<https://doi.org/10.17226/23395>
- Nefsky, Julia. 2018. Consumer Choice and Collective Impact. *Oxford Handbook of Food Ethics*. Oxford University Press.
- O'Neill, Jim et al. 2015. Antimicrobials in Agriculture and the Environment: Reducing Unnecessary Use and Waste. *The Review on Antimicrobial Resistance*.
<http://amr-review.org/>.
- O'Neill, Jim et al. 2016. Tackling Drug-Resistant Infections Globally: Final Report and

- Recommendations. *The Review on Antimicrobial Resistance*. <http://amr-review.org/>.
- Robbins, JA *et al.* 2016. Farm Size and Animal Welfare. *Journal of Animal Science*, 94(12): 5439-5455
- Sender, Ron *et al.* 2016. Revised Estimates for the Number of Human and Bacterial Cells in the Body. *PLOS Biology*. <https://doi.org/10.1371/journal.pbio.1002533>
- Sinnott-Armstrong, Walter. 2005. It's Not My Fault: Global Warming and Individual Moral Obligations. *Perspectives on Climate Change: Science, Economics, Politics, Ethics Advances in the Economics of Environmental Research*, Volume 5, 293–315.
- Spellberg, Brad *et al.* 2016. Antibiotic Resistance in Humans and Animals. Discussion paper for the National Academy of Medicine. Accessed July, 2017. <http://www.nam.edu/antibiotic-resistance-in-humans-and-animals>
- Trivers, Robert. 1971. The Evolution of Reciprocal Altruism. *The Quarterly Review of Biology* 146(1): 35-57.
- Van Boeckel, TP *et al.* 2015. Global Trends in Antimicrobial Use in Food Animals. *Proceedings of the National Academy of Sciences*, 112(18): 5649-54.
- Velasquez-Manoff, Moises. 2013. *An Epidemic of Absence: A new way of understanding allergies and autoimmune diseases*, (New York: Scribner)