

QUESTIONS

1. What are the main objections that McPherson considers to each premise of the argument in the model essay? Explain the main point of each objection in your own words in one or two clear sentences.
2. What exactly does McPherson's Anti-Complicity Principle say? How does it differ from the principles outlined in the introduction to this chapter? Explain using clear examples that illustrate the differences. Can you think of reasons why someone might think that McPherson's principle does a better or a worse job than the principles discussed in the chapter introduction of capturing what we often think is intuitively wrong in cases where a person might be described as "complicit in wrongdoing"?
3. Why does McPherson favor adding the Anti-Complicity Principle to the argument? What are the alternatives he considers? Why do you think he favors Anti-Complicity over alternative principles? How do you expect he would argue against alternatives?
4. Can you think of clear cases that could be taken to be counterexamples to the Anti-Complicity Principle? How do you think McPherson would respond? How would you respond?

MARK BUDOLFSON

The Inefficacy Objection to Utilitarian Theories of the Ethics of the Marketplace

Budolfson argues that the inefficacy objection raises an important problem for utilitarian arguments in consumer ethics, including for Peter Singer's arguments for vegetarianism. He examines an "expected consequences" response to the inefficacy objection on behalf of such utilitarians, which claims that the expected effect of one individual's consumption on production is roughly equal to the average effect on production of everyone's actions of that type. Rudolfson argues that we should not expect this response to work given readily available knowledge about the workings of markets. He argues that in light of those facts about the actual workings of markets, we should expect conclusions opposite to those of Singer's to be true. (Budolfson is a philosopher at the University of Vermont.)

Mark Rudolfson, "The Inefficacy Objection to Utilitarian Theories of the Ethics of the Marketplace." Printed with the author's permission.

The signature ethical problem of the global consumer society is our responsibility for the unethical practices that lie behind the products we buy.

—Peter Singer

Americans would be healthier if they reduced their consumption of animal products, and any desirable level of animal consumption could be sustained without factory farming. Furthermore, eliminating factory farms would save billions of animals from intense suffering each year, improve public health, reduce food prices in lesser developed countries, reduce energy costs by reducing the price of biofuels, reduce air and water pollution, reduce the irrational consumption of scarce groundwater resources, reduce greenhouse gas emissions, reduce health-care costs and thereby reduce the national deficit, and allow family farmers to compete more effectively against agribusiness conglomerates.

In light of the preceding, and the fact that the U.S. government could adopt efficient regulations that would largely eliminate factory farming without causing any weighty negative side effects, the U.S. government should adopt such regulations. However, in the meantime factory farming continues. Does this mean that each American *as an individual person* is required to stop consuming animal products from factory farms, at least when s/he can do so without dramatic inconvenience?

At first glance, it might seem that the answer to this individual-level question is “yes,” just as the answer to the collective-level question of whether we should eliminate factory farms is “yes.” However, even if it is clear what solution is best at the level of public policy, this does not mean that it is clear what *individuals* have ethical reason to do in the absence of such policy—especially if factory farming is a problem to be solved only by policy and cannot realistically be solved by individual ethics or by individual people “voting with their dollars.”

In what follows I argue that the question of what individuals are required to do about factory farms is indeed more complicated than it initially appears—and, for similar reasons, most questions about what individuals are required to do when products are

produced in morally objectionable ways are more complicated than they initially appear. As I argue at greater length elsewhere, I think these complications have general importance, because they ultimately call into question whether familiar ethical theories can explain many of the most important facts about modern moral life.

To see how things get complicated, in this paper I’ll focus on utilitarian arguments that it is wrong to consume animal products from factory farms, such as those offered by Peter Singer. The intuitive idea behind Singer’s arguments is that purchasing and eating meat from factory farms is wrong because it has unacceptable consequences on balance for welfare. For example, if I purchase and eat a factory-farmed chicken, the idea is that my gustatory pleasure is greatly outweighed by the suffering that the chicken experiences in order to bring me that pleasure; as a result, Singer would claim that the (expected) effect on overall wellbeing of my eating that chicken is unacceptably negative on balance, even if I really enjoy eating it—and similar reasoning explains why it is in general impermissible for individual people to consume animal products from factory farms and other products that are produced in a way that is wrong.¹

Furthermore, note that Singer’s reasoning can seem decisive even if one endorses an ethical view other than act utilitarianism. That is because on any plausible view, the welfare effects of one’s actions are ethically important. In particular, even on deontological and virtue theories, it is generally wrong to bring about significant harm or suffering for no significant reason, even if doing so would not violate any rights or agreements. For example, if I know that bad weather has destroyed the food supply for animals in my area except for some berries that grow deep in the woods on my property, it would be wrong for me to burn all of those berries in my outdoor fireplace simply to provide a very fleeting pleasant aroma, if I know that this would cause all of the animals in my area to starve to death painfully—and one doesn’t have to be a traditional consequentialist to endorse this thought. As a result, if Singer’s argument about eating meat are sound on the assumption that act utilitarianism is true, then it seems like they

are also sound on the assumption that any other plausible ethical theory is true, because if Singer is right about the welfare effects of eating meat, those considerations would trump any other considerations that a plausible ethical theory could say are relevant to determining what we are required to do.

However, even if we agree with Singer's premises about the magnitude of animal suffering in the world, and the comparative unimportance of gustatory and other human pleasures, there is an important objection to his argument—the inefficacy objection—that claims that his conclusion about the welfare effects of consumption by an individual does not follow, and, upon careful reflection, turns out to be false. That is because, according to the inefficacy objection, an individual's decision to consume animal products cannot really be expected to have any effect on the number of animals that suffer or the extent of that suffering, given the actual nature of the long and complex supply chain that stands in between individual consumption decisions and production decisions; at the same time, an individual's decision to consume animal products does have a positive effect for that individual. As a result, according to the inefficacy objection utilitarian premises about the value of suffering vs. pleasure, together with the actual empirical facts about the workings of the marketplace, entail that the (expected) welfare effect of an individual's decision to consume many ethically objectionable products is *positive* on balance, in contrast to what utilitarians like Singer assume. So, if the inefficacy objection is sound, it undermines the idea that it is generally wrong for individuals to consume products that are, on utilitarian grounds, produced in a way that is extremely wrong.

In light of this challenge from the inefficacy objection, it is perhaps no surprise that utilitarians have unified behind a particular response, which claims that the inefficacy objection invariably overlooks the significance of low-probability threshold effects. Here is Alastair Norcross's presentation of such a response:

Suppose that there are 250 million chicken eaters in the US, and that each one consumes, on average, 25 chickens per year . . . Clearly, if only one of those chicken eaters gave up eating chicken, the industry would not respond. Equally clearly, if they all gave up eating chicken, billions of chickens (approximately 6.25

billion per year) would not be bred, tortured, and killed. But there must also be some number of consumers, far short of 250 million, whose renunciation of chicken would cause the industry to reduce the number of chickens bred in factory farms. The industry may not be able to respond to each individual's behavior, but it must respond to the behavior of fairly large numbers. Suppose that the industry is sensitive to a reduction in demand for chicken equivalent to 10,000 people becoming vegetarians. (This seems like a reasonable guess, but I have no idea what the actual numbers are, nor is it important.) For each group of 10,000 who give up chicken, a quarter of a million fewer chickens are bred per year. It appears, then, that if you give up eating chicken, you have only a one in ten thousand chance of making any difference to the lives of chickens, unless it is certain that fewer than 10,000 people will ever give up eating chicken, in which case you have no chance. Isn't a one in ten thousand chance small enough to render your continued consumption of chicken blameless? Not at all. While the chance that your behavior is harmful may be small, the harm that is risked is enormous.²

The goal of this response from Norcross is to argue that in the cases at issue, for any single individual, the *expected effect* of that individual's action—in the particular example discussed here, the action of purchasing one chicken—is equal to the *average effect* of all actual acts of that type across society, in the sense of “expected effect” familiar from the notion of “expectation” in expected utility theory³—and similarly for other collective action situations in which the inefficacy objection might initially seem to show that utilitarianism has absurd implications by the lights of its own defenders such as Singer and Norcross. If this response were correct, then utilitarians would be able to respond to the inefficacy objection in a way that appears fully satisfactory by their own lights.

In what follows, I will argue that this response to the inefficacy objection presented by Norcross does not ultimately work, and that the inefficacy objection is ultimately sound.

To begin to see the problem with the response outlined by Norcross, consider the following case:

Richard makes paper T-shirts in his basement that say “HOORAY FOR CONSEQUENTIALISM!,” which he then sells online. The T-shirts are incredibly cheap to produce and very profitable to sell and Richard doesn't care about waste per se, and so he produces

more T-shirts than he is likely to need, and then sells whatever the excess turns out to be at the end of the month at a nearly break-even amount to his hippie neighbor, who burns them in his wood-burning stove. For many years Richard has always sold between 13,000 and 16,000 T-shirts each month, and he's always printed 20,000 T-shirts at the beginning of each month. Nonetheless, there is a *conceivable* increase in sales that would cause him to produce more T-shirts—in particular, if he sells over 19,000 this month, he'll produce 25,000 T-shirts at the beginning of next month; otherwise he'll produce 20,000 like he always does. So, the system is genuinely sensitive to a precise tipping point—in particular, the difference between 19,000 purchases and the “magic number” of 19,001.

Suppose that a consumer knows all of these facts about Richard's business, and is considering buying a T-shirt for himself. What is the expected effect on the number of T-shirts produced of that consumer purchasing a T-shirt? The correct answer is essentially zero, because given what is known about the history of demand for Richard's T-shirts and how production quantities are determined, there is virtually no chance that exactly 19,001 people are going to buy Richard's T-shirts this month and trigger a dramatic threshold effect—which, of course, is not to claim that there is *zero* chance of that happening, but rather that the odds of that happening—of exactly 19,001 of Richard's T-shirts being sold and no more—is certainly dramatically lower than 1/5,000 or any other number that would drive the expected effect of an individual buying one T-shirt anywhere near the consequence that 1 additional T-shirt is produced.

This shows that the reasoning behind Norcross's response is flawed, because insofar as his response is taken to show that consuming meat should be expected to have significant bad effects for animal welfare (i.e., equal to the average effect of each such purchase across all of society), similar reasoning would show that buying one T-shirt in the story above should be expected to result in approximately 1 additional T-shirt being produced, which is the wrong result.⁴ The problem with the reasoning is that it overlooks the fact that we can know enough about the supply chains in both cases to know that threshold effects are not sufficiently likely and are not of sufficient magnitude to drive the expected effect of

consumption anywhere close to the average effect. For the sake of giving a name to the more general mistake illustrated by Norcross's reasoning of illegitimately assuming that the expected effect of a single individual's action must equal the average effect of all actual actions of that type, we might call that mistake the “Average Effects Fallacy.”

In response, it might be insisted that there is a crucial disanalogy between the T-shirt case just described and e.g. our actual situation with respect to factory farmed animal products. One obvious difference might seem to be the amount of waste: in particular, in the T-shirt case, a significant amount of the product is “wasted.” However, although the T-shirt case is indeed a dramatization aimed primarily at making vivid why the reasoning behind the response is invalid as presented, upon further reflection there is less of a difference to a typical case involving the consumption of animal products than it might initially appear, and more importantly there is no crucial disanalogy with respect to the expected effect of an individual's consumption decisions. For example, consider the meat that goes out of date in a wholesaler's meat locker or on a supermarket shelf, and is then sold to a dog food plant or “rendered” into feed for other animals. Is that meat “wasted”? What is the difference between that meat and the “wasted” T-shirts in the story above that Richard sells to his neighbor?

For current purposes it doesn't really matter whether such things are labeled as “wasted” or not—what matters is that there is reliably at least a small amount of “wasted”-like meat at each stage of the supply chain that serves as sufficient “slack” to create buffers that prevent an individual's decision to purchase meat from making any difference to the number of animals that are produced at the far other end of the supply chain. Most importantly, the crucial issue is not about the *magnitude* of these buffers, but rather about their *reliability*: as long as we can know—as we can—that there are sure to be buffers of non-trivial size throughout the supply chain (even if they are not nearly as large as the buffers in the T-shirt example), that reduces the probability of a single individual making a difference to a level that appears to become nearly infinitesimal.

So, the key empirical argument here relevant to consumer ethics is that (1) Many products we consume

are delivered by a massive and complex supply chain in which there is some *reliable* amount of *slack*—in the form of waste, inefficiency, and so on—at many links in that chain. (2) That slack serves as a buffer to absorb any would-be effects from the links before. Furthermore, (3) Production decisions are insensitive to the informational signal generated by a single consumer because the sort of slack just described together with other kinds of noise in the extended transmission chain from consumers to producers, and are also insensitive to a single individual's action because production decisions themselves are generally not made in the way assumed by utilitarians in their responses to the inefficacy objection. (Conclusion 1) This ensures that significant-enough threshold effects are not likely-enough to arise from an individual's consumption decisions to justify equating the expected effect on wellbeing of an individual's decision with anything approaching the average effect of such decisions. As a result, (Conclusion 2) For many products in modern society, it is empirically implausible that even a lifetime of consumption decisions by a single individual would make any difference to quantity produced and thus the harm that lies behind those products. So, (Main Conclusion) Utilitarianism cannot explain why it is wrong to consume those products.

For an illustration, consider the supply chain for American beef. When ranchers who own their own grazing land decide how many cattle to raise, their decisions are sensitive to their own financial situation, the number of cattle their land can support, the expected price of any additional feed that will be needed, bull semen and other “raw materials” that go into cattle production, and the expected price that the cattle will fetch when they are ultimately sold to feedlots. Of these, small changes in the last item—the price that cattle will fetch at the feedlot—are of the least importance, because insofar as ranchers judge that capital should be invested in raising cattle rather than other investments, they will tend to raise as many cattle as they can afford to breed and feed within that budget, letting the ultimate extent of their profits fall where it may at the feedlot. Many ranchers also use the nutritional well-being of their herd as a buffer to absorb adverse changes in market conditions, feeding their cattle less and less to whatever

point maximizes the new expectation of profits as adverse conditions develop, or even sending the entire herd to premature slaughter if, say, feed prices rise to levels that are unacceptably high. This serves to shift the ranchers' emphasis in decision-making relevant to herd size even further away from the price of beef. As a result, even if an individual's consumption decisions managed to have a \$0.01 effect on the price of cattle at feedlots, the effect on the number of cattle produced would be much smaller than it would have to be in order for the (vanishingly small) probability of such a threshold effect to justify equating the expected marginal effect of an individual's consumption of beef with the average effect of such consumption decisions. More importantly, because animal production is so many links in the supply chain away from grocery stores and restaurants, and because the intervening links typically involve some small but non-negligible amount of waste, inefficiency, and other forms of slack serve as a buffer to absorb any effect that your personal consumption might otherwise have. These facts suggest that there is good empirical reason to think that the expectation of a single individual's consumption decisions on production is nearly zero and is not to be equated with the average effect of similar consumption decisions across society, contrary to what the response to the inefficacy objection by Norcross suggests. A similar upshot emerges even in a more vertically integrated industry such as the poultry industry, where demand is relatively inelastic; nonetheless, because profits are dependent mostly on the cost of inputs such as feed and fuel, a similar story about the expected effect of a single consumer's purchases remains true.⁵

To more completely evaluate the prospects for utilitarian arguments regarding consumer ethics, another important consideration is that even if you would convince many others to be a vegetarian by becoming one yourself, that does not translate into strong welfare-based reasons to become a vegetarian. That is because even if your vegetarian lifestyle ultimately caused, implausibly, one hundred others to become vegetarians who would not otherwise have done so, their collective consumption decisions might still not have any appreciable effect on the number of animals that are raised and mistreated,

because the actual mechanisms in the marketplace may be insensitive to the distributed effects of even one hundred consumers. Of course, this reasoning does not hold true when applied to an influential person like Peter Singer who really does influence enough people to make a difference, but it does hold true when applied to almost everyone else, which means that utilitarianism does not require most individuals to become vegetarians, even if it requires a few influential people like Peter Singer to be vegetarians. For example, just as morality does not require each of us to act as if we had the talents, influence, and resources that Warren Buffett has, so too morality does not require us to act as if we had the talents, influence, and resources that Peter Singer has.

A related complication is that individual vegetarian acts often have negative unintended consequences that must also be properly accounted for. Like the positive “indirect effects” discussed in the preceding paragraph, this is also a genuine complication that must be taken into proper account in a consequentialist analysis. For example, if I am a vegetarian, I might easily alienate others with my vegetarian acts insofar as they tend to be interpreted as self-righteous, and thus cause others to adopt a policy of never reducing their consumption of meat and never taking vegetarian arguments seriously. And if vegetarians are generally interpreted as self-righteous, that might lead to a consensus among most members of society that vegetarians are radical, self-righteous jerks who should not be taken seriously and who should be scoffed at by others—which then raises the cost of making vegetarian choices for everyone, and is counterproductive in other ways.⁶ (Here it may be useful to note that many undergraduate students respond to vegetarian eaters by pledging to eat more meat to “cancel out” the effects of the vegetarians.)

In sum, the key issues here relevant to consumer ethics are that many products we consume are delivered by a massive and complex supply chain in which there is waste, inefficiency, and other forms of slack at each link. Arguably, that slack serves as a buffer to absorb any would-be effects from the links before. Furthermore, production decisions are insensitive to the informational signal generated by a single consumer because the sort of slack just described

together with other kinds of noise in the extended transmission chain from consumers to producers ensures that significant-enough threshold effects are not likely enough to arise from an individual’s consumption decisions to justify equating the effect of an individual’s decision with anything approaching the average effect of such decisions. As a result, for many products in modern society, it is implausible that even a lifetime of consumption decisions by a single individual can be expected to make any difference to the harm that lies behind those products, and so it is implausible that we have strong utilitarian reasons not to consume those products.

A more general philosophical upshot is that when evaluating utilitarian reasons for individual action, the knowledge available to agents about the mechanisms at play in their specific situations matters greatly, especially when outcomes depend on collective action—and it is a mistake to think that there is any simple argument that shows that individuals always have decisive utilitarian reasons to “cooperate” and perform “pro-social” actions whenever the stakes for society are high of our collective action.

NOTES

1. See Peter Singer (1980). “Utilitarianism and Vegetarianism,” *Philosophy & Public Affairs*.

2. Alastair Norcross (2004). “Puppies, Pigs, and People,” *Philosophical Perspectives*, pp. 232–3. See also Peter Singer, *ibid.*, and Shelly Kagan (2011). “Do I Make a Difference?” *Philosophy & Public Affairs*.

3. In particular, the “expected effect” of an action in this sense is the *expectation* associated with that action, based on the sum of all of the values of all of the possible outcomes of that action weighted by their probability conditional on that action. See Michael Resnik (1987), *Choices: An Introduction to Decision Theory*, U of Minnesota Press for further discussion of expected utility theory.

4. Note that this result is not undermined by the observation (often enthusiastically made by consequentialists) that in collective action situations, as the probability of making a difference goes down, the magnitude of the difference that would be made goes up. As the discussion above illustrates, what matters is whether the difference that would be made increases in a way that is relevantly proportional to the decrease in the probability of making a difference—and as the discussion here illustrates, in real

world collective action problems it is often empirically unrealistic to think that it does.

5. See for example the comments of poultry industry expert Ed Fryar in Ryssdal, K. (2015). “Why chicken wings cost more this time of year,” *Marketplace*, 5 January 2015: <http://www.marketplace.org/topics/business/why-chicken-wings-cost-more-time-year>

6. For some other suggestive empirical results that vindicate this as a serious general worry, see Julia Minson and Benoit Monin (2011), “Do-Gooder Derogation,” *Social Psychological and Personality Science*, and Benoit Monin et al. (2008), “The Rejection of Moral Rebels: Resenting Those Who Do the Right Thing,” *Journal of Personality and Social Psychology*.

QUESTIONS

1. Does Budolfson consider the possibility that by deciding to be a vegetarian one might influence many other people, leading them to become vegetarians, who might then influence other people, and so on? (What does Peter Singer say about this in “Utilitarianism and Vegetarianism”?) Does Budolfson miss something important here, or are his arguments convincing?
2. If Budolfson’s arguments are correct, does it follow that it is permissible to consume everything? Why or why not?
3. Do Budolfson’s arguments assume that utilitarianism or some form of consequentialism is true? Explain.
4. How might someone argue that the inefficacy objection also raises problems for deontological ethical theories? Give the best example possible of how someone might argue that they do not raise problems, and of how someone might argue that they do raise problems.
5. How would Budolfson’s argument apply to other cases? What would be the implications for what individual people have reason to do in the cases discussed by Williams and Hill, cases of voting in elections, cases (discussed by papers in the Further Reading section) of divesting from investments in institutions doing wrong discussed by Appiah, and cases discussed by Kagan and Nefsky, such as climate change? Does utilitarianism or Budolfson’s arguments gain or lose plausibility from careful consideration of these cases and related arguments? Explain.

ELIOT MICHAELSON

Act Consequentialism and Inefficacy

Michaelson argues that even if Budolfson’s argument is convincing against an “expected consequences” response to the inefficacy objection, this may not undermine the force of consequentialist considerations in our actual world, and so Budolfson might be incorrect about what consequentialism *actually* implies in the consumer ethics cases we care about.

Eliot Michaelson, “Act Consequentialism and Inefficacy.” Printed with the author’s permission.