

Act Consequentialism and Inefficacy  
Penultimate Draft  
Eliot Michaelson

A variety of purchasing and eating practices—vegetarianism, veganism, pescatarianism, locavorism, etc.—are often motivated along broadly act consequentialist lines. That is, these sorts of practices are often motivated by the thought that individual decisions not to purchase meat (I'll focus on just this example here) *matter* for the world. Or, at the very least, that this sort of decision has a decent chance of *matter*ing. Fewer animals will lead short and unhappy lives, fewer animals will die badly, and the world will generally be a better, happier place. The reasoning here seems hard to dispute: you should do what maximizes expected goodness. Purchases *signal* the market to produce more or less of a certain good, or to deliver it in one way as opposed to another. These signals can be expected to result in better or worse outcomes. You should therefore signal the market to produce fewer goods that, overall, reduce the goodness of the world. So unless one wants to cast doubt on the entire apparatus of market signaling, it will be hard to dispute that—assuming that treating animals how we do on contemporary factory farms is generally a bad thing—we will have strong consequentialist reasons to refrain from purchasing meat. Or so it would seem.

There is, however, a serious concern with this line of reasoning, and one that has been lurking in the background of such arguments for some time. Suppose that one derives some non-trivial enjoyment from eating meat. How should the value of this enjoyment be weighed against the value of this market signaling effect? The enjoyment, recall, is guaranteed. The signal may or may not lead to a change in production and distribution practices. And, at the end of the day, consequentialists care about *consequences* (or at least expected consequences). It looks, therefore, like empirical facts about *how likely* it is that one's purchasing habits will have an effect on subsequent production and distribution practices may matter quite a bit for whether act consequentialism will recommend adopting a vegetarian set of eating and purchasing practices.

To illustrate the point, consider the following example: suppose that all of the world's chickens are produced by a global monopoly called "Chicken Demon".<sup>1</sup> Chicken Demon is run by an eccentric billionaire, "The Colonel", who loves nothing more than torturing and killing chickens. Chicken Demon raises as many chickens in as many countries the world over as the available resources will permit; the operation is not sensitive to profits and The Colonel effectively has an endless supply of money to support Chicken Demon if it slides into unprofitability. The Colonel is a good free-market capitalist, however, so if people are willing to pay for the byproduct of his chicken-torturing and -killing enterprise, i.e. chicken carcasses, he is more than willing to sell these carcasses for a tidy profit. The important thing is that The Colonel, and hence Chicken Demon, isn't motivated by the amount of money coming in. What The Colonel cares about is torturing and killing chickens, and he has the resources to pursue that goal to his heart's content. So while *there is* a market signal generated by global chicken-demand, Chicken Demon's behavior is *insensitive* to that signal. In this scenario, act consequentialism can offer no reason to refrain from eating chicken. One's choosing to eat chicken or not to do so will send a market signal, but that signal is *guaranteed* to have no effect on the morally relevant behavior, i.e. torturing and killing chickens. So the act consequentialist's argument for vegetarianism turns out to be *contingent* on certain empirical facts—and, in particular, on producers exhibiting a certain

---

<sup>1</sup> This is an exaggerated version of Mark Budolfson's "waste" example from his "The Inefficacy Objection and the Problem with the Singer/Norcross/Kagan Response".

degree of sensitivity to market signaling. For act consequentialism to recommend being vegetarian, there needs to be at least some chance that one is going to make a difference.

This much has been acknowledged by consequentialists like Peter Singer, Alastair Norcross, and Shelly Kagan. What these authors have claimed is that our world is relevantly different from the world of my Chicken Demon example. In our world, the individuals and corporations that produce chickens and other foodstuffs are sensitive to the decisions of consumers. The question is just how sensitive they are, and how we are to conceive of this sensitivity as mattering for the subsequent state of the world. To help us get this question in view, let us consider a contrasting scenario, one of perfect market sensitivity.

Suppose once more that there is a global chicken monopoly, “Chicken Angel”. Chicken Angel, run by the eccentric billionaire “The Corporal”, is concerned to eliminate all waste and unnecessary harm in the chicken-production industry. As a result, when one wants to purchase a chicken, one enters an order at a Chicken Angel franchise. Chicken Angel then undertakes to raise one more chicken in their facilities. One to two months later, a chicken carcass is delivered to one’s door (in the unlikely event that one’s chicken dies of natural causes before maturity, delays may ensue). In contrast to Chicken Demon, Chicken Angel is *perfectly sensitive* to market signals; for every purchase of a chicken, one more chicken is raised and killed. Thus, supposing that the pleasure one derives from eating that chicken is less weighty, morally speaking, than the harm done to the chicken, act consequentialism will tell straightforwardly against ever ordering a chicken from Chicken Angel.<sup>2</sup>

Of course, we live in neither the world of Chicken Demon nor the world of Chicken Angel. And we might well wonder: what does act consequentialism recommend in our world? Singer, Norcross, and Kagan have all claimed that, in the morally relevant sense, our world is more like the world of Chicken Angel. We can expect to make a difference. Thus, we should refrain from purchasing meat. Recently, Mark Budolfson (“You Don’t Make a Difference”) has argued that the reasoning behind Singer, Norcross, and Kagan’s confidence here is fallacious, and based on what he calls the “Average Effects Fallacy”. I will argue, briefly, that neither is quite right.

Singer, Norcross, and Kagan all reason more or less as follows. Suppose that chickens are in fact purchased in batches of 25. And suppose that your local butcher will order a new batch whenever 23 chickens have been sold. At the end of the day, he throws out any remaining chickens. This means that you will be a *difference maker* only if you purchase the 23rd chicken. Otherwise, nothing will change in the world as the result of your purchase (the butcher will never go out of business as the result of selling one less chicken). Your chance of being a difference maker would seem, naturally enough, to be 1 in 25. So it looks like you have a 1 in 25 chance of being a difference maker with respect to 25 chickens. The expected badness of your act is thus equivalent to killing one chicken yourself.<sup>3</sup> In other words, the badness of the purchase is

---

<sup>2</sup> Two notes here: first, it may well be that there are restaurants in the real world that are relevantly similar to Chicken Angel, though on a smaller scale. Second, something akin Non-Identity Problem (see chapter XX) rears its head here. Basically, we might ask whether it is better for one more chicken to live a bad life than to lead no life at all. I will tentatively assume here that the answer is “No,” as reasoning along these lines leads to absurd consequences.

<sup>3</sup> Cf. Singer 1980, Norcross 2004, and Kagan 2011.

equivalent to the average badness of twenty-five people choosing to kill twenty-five chickens together. This is what Budolfson means when he says that Singer, Norcross, and Kagan equate the moral value of an individual's meat purchase with the average effect of everybody's purchases.

The problem, as Budolfson points out, is that purchases often aren't distributed in a random series—and, in fact, we generally know that they are not. Suppose that, aside from your own purchases, the number of chickens purchased per day at the local butcher can be accurately modeled by rolling two twelve-sided dice. There are days when just 2 chickens are sold, and days when 24 are. But there will be many more days when 12 chickens are sold. Likewise, there are many, many days when fewer than 12 chickens will have been sold by the time you find yourself considering whether to buy a chicken than there will be days when 22 have already been sold. Given this sort of purchasing series, the expected badness of your chicken purchase will be far less than the badness of killing one chicken. Days when you would make the difference—tipping the shop from 22 to 23 chickens—will be exceedingly rare. Supposing that you are aware of all this, then you should expect to have far less of a bad effect on the world by purchasing a chicken at the butcher shop than you should expect to have if you were to kill a chicken yourself. That, in turn, makes it substantially easier for the badness of this act to be swamped by its potential upside, e.g. your gustatory pleasure. That is, to justify buying a chicken, the pleasure one will obtain from eating that chicken needn't outweigh the badness of killing a chicken, but rather the badness of killing a chicken multiplied by the slight chance that one's action will actually lead to a chicken getting killed.

But is the badness of purchasing a chicken *in fact* swamped in our world? Here, things are far from clear. On the one hand, reflection on how the meat-production industry operates indicates that, in the actual world, meat producers may be relatively insensitive to market signals. For one thing, supply chains involve a number of layers (retailers, distributors, wholesalers, slaughterhouses, producers, etc.), offering multiple opportunities for an individual market signal to get lost along the way. In other words, for one's signal to make a difference, it must get all the way from the point of purchase up to the farmers who are actually raising the relevant sort of animal. Every step along the way, however, is at least somewhat waste-tolerant, meaning that many, many signals will get lost before one actually gets through. Each layer is like the butcher we considered above, who is only sensitive to whether 23 chickens have been sold. Another relevant consideration is that demand for meat products tends not to fluctuate randomly in the real world, aside from chance events like disease scares. Rather, demand for meat grows roughly in accord with population and wealth. What's more, meat distributors may also be willing to supplement standing orders (for, say, 25 chickens) with much smaller batches (of, say, 5 chickens) when demand unexpectedly spikes in a region. Finally, there is excess demand for meat products in the real world—particularly in the developing world—meaning that the question is not so much whether so much meat will be produced but *just how profitable* it will be to sell it. Given how cheap it is to produce meat in many places around the world, what stands in the way of more meat production is plausibly the lack of available land and water rather than a lack of potential profit. We thus appear to be far closer to the world of Chicken Demon than we might initially have anticipated.

On the other hand, the negative effects of raising and killing chickens are also far worse in our world than we have been supposing so far. Labor exploitation is rampant in the meat-production industry, and is not easily avoided by purchasing free-range or other “ethically-treated” meat products. Consumers have shown relatively little interest in treating workers, as opposed to an-

imals themselves, more humanely. Likewise, the negative environmental effects of raising chickens in the way we do must also be factored in.

Where does this leave us? Do we make a difference via our consumption habits? The answer would seem to be: possibly, but very little if we do. Still, that small chance may still be enough to offer a consequentialist justification for being vegetarian. One might, for instance, think that it is unacceptable to give one's child some medicine that will make her recover faster from a minor cold if that medicine has even a one in a million chance to killing her. On the other hand, we might find this trade-off acceptable if what we're dealing with isn't a minor cold, but something causing her significant discomfort. It seems that much hinges on the costs that are to be tallied on the other side of the ledger. Many people think of giving up meat as a serious cost in terms of gustatory pleasure, but that disvalue is likely to be front-loaded. Once one learns to cook vegetarian and navigate the vegetarian culinary world, there is scant evidence that one ends up enjoying food any less, *pace* the claims of many a committed meat-eater. What's more, the potential health benefits of vegetarianism are well-documented.<sup>4</sup> These considerations make it unclear just how the ledger will lean; much depends on some actual, and rather complicated, empirical facts that we cannot simply discern from the armchair.

What Budolfson's response to the Singer/Norcross/Kagan line should make clear, however, is the following. Whereas act consequentialism might have initially seemed appealing in that it explained why, if we ought collectively to refrain from purchasing meat, then we ought individually to do so, this transition from the collective to individual ought relies on some highly questionable empirical assumptions. The act consequentialist can only justifiably endorse this sort of argument if certain contingent features of the world turn out to be a particular sort of way. What's more, it is highly questionable whether the world turns out to be this way. For dedicated vegetarians, this is likely to seem like an odd result. What was wanted was an argument for vegetarianism *per se*, not an argument for vegetarianism that is held hostage by certain recondite details about the market for meat products—let alone an argument that depends on certain recondite details that look unlikely to obtain in the real world! If one is looking for an argument for vegetarianism *per se*, then Budolfson is surely correct in pointing out that act consequentialism looks poorly situated to provide such an argument.

## Bibliography

Key, Timothy J., Paul N. Appleby and Magdalena S. Rosell, "Health Effects of Vegetarian and Vegan Diets," *Proceedings of the Nutrition Society* 65(1): 35-41.

McEvoy, Claire T., Norman Temple and Jayne V Woodside, "Vegetarian Diets, Low-Meat Diets, and Health: A Review," *Public Health Nutrition* 15(12): 2287-2294.

---

<sup>4</sup> Cf. Key et al. (2006) and McEvoy et al. (2012).