The Impotence of the Causal Impotence Objection

Alastair Norcross
University of Colorado Boulder

Abstract: Many significant harms, such as the mass suffering of animals on factory farms, can only be prevented, or at least lessened, by the collective action of thousands, or in some cases millions, of individual agents. In the face of this, it can seem as if individuals are powerless to make a difference, and thus that they lack reasons, at least from the consequentialist perspective, to refrain from eating meat. This has become known as the "causal impotence" problem. The standard response is to appeal to expected utility calculations. Recently, this response has been attacked, mostly on the grounds that the relevant causal mechanisms are more complex than its proponents are said to assume. In this paper, I argue that the attacks are unsuccessful, both at undermining specific expected utility calculations urged by me and Kagan, or even at showing that significantly different expected utility calculations wouldn't justify the relevant behavior.

In recent years, a number of philosophers have urged the causal impotence objection against expected utility arguments for individual behavior aimed at addressing large-scale harms, such as the arguments given by Norcross (2004) and Kagan (2011) against eating meat.[1] For example, perhaps there are situations we can find ourselves in, where we know, *with certainty* (not just a very high degree of subjective probability) that our purchase and/or consumption of a piece of meat will make no difference to the number of animals raised for human consumption. John Harris and Richard Galvin (2012) present such a case, involving the last chicken sandwich at a Wendy's in Kingman Arizona at 3AM, seconds before closing time. Perhaps we can know enough about exactly what will happen if we buy the sandwich in this specific case, versus what will happen if the sandwich is thrown away, to reduce our subjective credence in our act of buying having any effect on how many animals are raised to zero. Not merely very close to zero, but literally zero. This would involve knowing, not just either that the restaurant would not record a difference in their official report to whomever they report to, or that such a report would have, again, literally zero chance of making a difference at any other stage in any causal process that might lead, however indirectly, to differences being made in decisions about numbers of animals to raise, but also that our

choice of purchasing and consuming this chicken sandwich would have no chance of affecting our own future behavior, or that of others, in any way that could be part of a causal chain, no matter how indirect, leading to such aforesaid differences in animal raising decisions. With all these specifications, it seems pretty unlikely that we will ever find ourselves in such a situation. I'm pretty sure I've never been in such a situation myself. Nevertheless, it's certainly possible to find oneself in such a situation. Notice that the relevant situation is not one in which, as a matter of fact, your behavior will make no difference to the relevant outcome (in this case, the number of animals being raised for human consumption). It is rather a situation in which you are justified in believing that there is a zero probability of your behavior making such a difference. Not a very small, but nonzero probability, but literally a zero probability. So, what should we say about such situations? Well, in such rare situations, you really would have no animal-welfare related subjective reason to abstain from the relevant behavior. But those are not the situations that the vast majority of meat eaters find themselves in almost all the time.

Although Harris and Galvin produce an example in which we are supposed to think that we are justified in believing that there is a zero probability of our behavior having any relevant effect, their main strategy in response to me is to stress how unlikely we should consider the possibility of such effects to be. Without any attempt at measuring the probabilities, though, either subjective or objective, this response simply fails to engage with the expected utility approach. My schematic example assumes a probability of one in ten thousand, or one hundredth of one percent. That's certainly, by any normal understanding of the term, a very small chance.

Others have pointed out that Kagan's and my approach seems to assume a simple linear relationship between demand, as measured by consumer consumption behavior, and supply, as measured by the total number of animals of various kinds raised in various ways. The actual market, though, is more complicated. As a representative example, consider Julia Nefsky:

> Some stores may have more complex ordering strategies, in which decisions depend on a variety of factors, including upcoming marketing plans, the availability and cost of other similar products, news about one's competitors, and a wide range of statistics about past sales [...]. The factory farm most likely chugs along producing as much as it can, given the space and resources it has. If one particular butcher orders fewer

> chickens than usual, this might result in the distributor searching for a new client, rather than in a decrease in production. The factors that go into deciding whether to increase or decrease future production, and by how much, might include not only facts about how the distributors have done in their sales, but also the physical limitations of the space, the financial burdens of expanding, and so on. (2012)[2]

None of the critics suggest that consumer demand is *irrelevant* to supply, which would be really silly, but rather that there are other factors which help determine supply in various circumstances. But presumably they would all admit that, in some obvious sense, demand is the *fundamental* factor, at least in anything like a market economy. After all, it really doesn't matter how many buffers you have built in to absorb waste or inefficiency, or how many limited-time price adjustments are being offered, or what tax incentives or other legal features are in place, if demand is low enough, no factory farm is going to chug along producing as much as it can. Assume otherwise. Picture the annual board of directors meeting at Perdue Farms chicken conglomerate. Frank Perdue: "How many broilers did we produce this year?" CFO: "2 billion, same as last year, and the year before, and the year before, and, in fact every year since we opened our last farm. Remember our motto, 'Chug Along As Usual.'" FP: "So, how was demand this year? How many people ate chicken, and how many of those bought how many of our broilers?" CFO: "Well Frank, demand is down again. It appears there are now fewer than one million chicken eaters in the US. We estimate that about two hundred and fifty thousand of them bought about ten million of our broilers." FP: "Hmm. So, for next year, what do you say, shall we chug along as usual? Another two billion?" CFO: "Yes, that seems right to me. As I learned in Business School, the surest way to success in a market-based economy is to pay no attention whatsoever to consumer behavior."

The reason why this is absurd is that consumer behavior, in the form of demand, is clearly relevant to supply, and, ultimately, more important than any other factor in anything like a market economy. Granted, say the critics, but we don't live in a world in which demand for chicken in the US is down to just a few million a year. In the US, over eight billion broiler chickens are raised each year on factory farms, and most of them are eventually sold and eaten. You can't simply take the number of chickens produced, and divide it by the number of chicken consumers to get the expected utility of one consumer's change in behavior making a difference. But there must still be thresholds, perhaps vague ones,

otherwise we could find ourselves in the absurd situation of the previous example. Other factors may well prevent the response to demand from being linear. Perhaps, when ten thousand fewer chicken eaters purchase a quarter of a million fewer chickens, a large part of the slack will be taken up by these other factors, so that production won't adjust at that point, but after even more consumers change their behavior. Perhaps it will take twenty or thirty thousand people changing their behavior to bring about a production adjustment equivalent to the consumption of only ten thousand. Or maybe, though this seems pretty unlikely, even forty or fifty thousand people would have to purchase one or one and a quarter million fewer chickens, in order to produce a reduction in production of a quarter of a million chickens, which could render the expected reduction in chicken breeding of each foregone chicken as low as one fifth of a chicken. Of course, if that happens, there will be less room for slack in the subsequent arrangement, suggesting that the next threshold will take fewer changes in behavior to reach, but perhaps still more than the subsequent adjustment would take into account. So, what is the relevance of these considerations to expected utility calculations? If we know with certainty that there won't be a change in behavior big enough to precipitate either a decrease or an increase in animal production, then the expected utility in terms of animal welfare of changing our diets is zero. But, of course, we don't know that. In fact, it's easy to see both that the numbers of animals being raised for food does in fact vary from year to year, and that the numbers of most animals raised for food, including chickens, is on a diminishing trend, compared with the total population (and surveys suggest that any increases in total meat consumption, are mostly driven by immigrants from countries where meat is more expensive, who, at least in the short term, increase their consumption of meat). But perhaps we can know with certainty that there will never be enough changes in behavior to produce a counterfactual difference in animal rearing corresponding to the numbers of animals no longer consumed. Perhaps we can know with certainty that the market will never, in fact, be completely efficient. The first thing to say in response to this is that we have absolutely no reason to think that we could know any such thing. Perhaps the next threshold that our behavior contributes to reaching (or to preventing being reached) won't represent a change corresponding to the change in numbers of animals consumed, but, unless we know with certainty that no other threshold will ever be reached, the story doesn't end there.

In fact, there is considerable evidence that markets are becoming *more* efficient, not less. This is mostly due to ever more sophisticated mechanisms for tracking and analyzing consumer behavior. Most

consumers pay for purchases with credit cards, which are tracked by sophisticated computer programs. My consumer behavior is analyzed, and used to tailor advertising on my computer directly to me.

But let's grant, for the sake of argument, that, because of the complications of an inefficient market, the expected counterfactual reduction in animal rearing of my giving up eating *n* animals of a certain sort is considerably smaller than *n*. How much smaller than *n* would it have to be, in order for the prudential benefit of my continuing to eat n animals to outweigh the harm of raising that smaller number? Well, first of all, there would have to *be* a net expected prudential benefit of eating *n* animals. Almost certainly, the net effect on my own welfare of eating animals is negative. Even if the gustatory pleasure really couldn't be equaled by alternative food sources (almost certainly false), the negative effects on my health would almost certainly outweigh any gain in pleasure. Then, I would have to consider any further harms that raising that number of animals would cause, beyond the suffering of the animals themselves. The most obvious such harms are environmental. By now it is well known that animal agribusiness is the largest single contributor to anthropogenic global warming. Further, there would be the harms to our character (and the resulting harms to sentient beings affected by our behavior) that are caused by the perpetuation of a system that regards animals as resources for us to exploit. And this, of course, assumes that we have already factored in the more indirect ways that our behavior might lead to a reduction in the numbers of animals reared, such as precedent and example-setting effects, and the kind of support and bolstering we can give to organizations that are working for change by legislative or other means. To the extent that our dietary behavior contributes to solidarity among those opposed to rearing animals for food, and to the extent that such solidarity helps to achieve positive change, as it undoubtedly does, our contribution to solidarity itself makes a change in animal rearing more likely. It is important to remember that "making a difference" is not just a simple matter of a direct causal link between our behavior and a change in the numbers of animals raised for food. There are many more indirect ways in which our behavior can make a difference. Once we see this, the idea that our dietary choices can't make a difference is revealed as absurd.

But, let's suppose that we have really taken all these considerations into account, and that we are confident that, in our case at least, there will be a net prudential benefit to eating meat. How small would the expected change in animal rearing have to be in order for the prudential benefit to outweigh it? Suppose that the expected reduction in chicken rearing of eating one chicken is only a twentieth of a chicken (far smaller than is

even remotely justified by anything the critics say). And suppose that the net prudential gain of eating the chicken is the entire pleasure of eating that chicken. Even though there is no nutritional gain, compared with readily available alternatives, suppose, *per impossibile*, that you would gain no gustatory pleasure whatsoever from eating an alternative. Now, consider the life of a factory chicken. Most commercial broiler chickens live for around six weeks. One twentieth of that is just over two days. So, even granting far more than anything the critics have said actually warrants, you would have to believe that your pleasure from eating that one chicken is more morally significant than the suffering that such a chicken endures during more than two whole days of its tortured existence. Think about that for a minute. *Really* think about it. What kind of moral monster could actually maintain this?

Given what we know about the ways the markets work, and especially about the ways in which sophisticated technology is getting increasingly good at tracking and analyzing consumer behavior, I suspect that we are epistemically justified in assigning something pretty close to the expected utility values suggested by me and Kagan. But even if we think that the complications in the market will reduce efficiency by a factor of twenty, and nothing any of the critics have said justifies anything remotely close to such a large degree of inefficiency, we would have to be moral monsters to think that our net prudential benefits would outweigh even the much smaller expected harms that are the outputs of the inefficient mechanisms.

So, why is the causal impotence objection still being urged? For some, it's probably a desperate desire to justify their continuing to eat meat, despite decisive reasons not to. For others, who already agree that they shouldn't eat meat, it's an even more desperate desire to find something wrong with consequentialism. Recall the absurd lengths to which this desire drives Philippa Foot (1983). She is even prepared to assert that there is no sense in which it is all-things-considered overall better that only one person dies than that millions of others die instead. But mostly, I suspect, it is a matter of paying insufficient attention to the difference between very small chances and zero chances. Humans are terrible at responding intuitively to large numbers whether as numerators or denominators. Recall the worries that were raised over the operation of the large hadron collider at CERN. A rumor went around the internet that the operation of the collider could cause the formation of a stable black hole, which would then destroy the entire planet (and maybe a lot more, I can't remember the exact details). I remember seeing a physicist interviewed on television about this very issue. The interviewer pressed him about this possibility. He responded that the worries were overblown,

because the chances of such a thing happening were very small. "How small?" asked the interviewer. He responded that the chances were no more than about one in ten million. I nearly fell off my chair. Don't they teach elementary arithmetic to physicists these days? Yes, one in ten million is, by most usual standards, a very small chance. But if what it is a chance of is the destruction of the whole planet, with more than seven billion human beings, and many billions of other sentient animals, and the subsequent prevention of the existence of who knows how many more, then one in ten million is nothing like small enough. Consider just the currently existing humans who would be killed by the destruction of the planet. Seven billion, multiplied by one in ten million, is seven hundred. If seven hundred people would definitely be killed by the operation of the collider, would that physicist still argue that we should do it? I have been assured by people who claim to know about these things that that physicist was, in fact, wildly off in his probability calculations. In fact, I've been told that there is really no reason to assign any finite probability greater than zero to the black hole possibility. The physicist was probably trying to convey that it was, in some sense, a real possibility (not logically or nomologically impossible), but that it was negligibly small, and groped for a probability that was close enough to zero, without actually being zero, as to make no difference. But of course, one in ten million is nowhere near close enough to zero to make no difference in this particular case, because of what is at stake. The destruction of the entire planet is not usually at stake in the kinds of cases that invite discussion of the causal impotence objection, such as climate change, and the large-scale suffering of animals raised for human consumption. Nonetheless, the potential (and increasingly actual) harms caused by global warming, and the massive actual harms caused to animals by agribusiness, are large enough to render even seemingly tiny probabilities ethically significant. The causal impotence objection is itself impotent in the face of these considerations.

Notes

[1] The causal impotence objection is also urged in the context of other behaviors, such as those aimed at combatting climate change, and even voting. Considerations of space limit my discussion in this paper to the context of eating meat.

[2] See also recent work by Mark Budolfson and Gary Chartier urging very similar considerations.

Works Cited

Foot, Philippa. (1983) Utilitarianism and the Virtues. *Proceedings and Addresses of the American Philosophical Association* 57(2): 273-283.

Harris, John and Richard Galvin. (2012) 'Pass the Cocoamone, Please': Causal Impotence, Opportunistic Vegetarianism and Act-Utilitarianism. *Ethics, Policy & Environment* 15(3): 368-383.

Kagan, Shelly. (2011) Do I Make a Difference? *Philosophy and Public Affairs* 39(2): 105-141.

Nefsky, Julia. (2012) Consequentialism and the Problem of Collective Harm: A Reply to Kagan. *Philosophy and Public Affairs* 39(4): 364-95.

Norcross, Alastair. (2004) Puppies, Pigs, and People: Eating Meat and Marginal cases. *Philosophical Perspectives* 18: 229-245.