Counting Your Chickens

Yoaav Isaacs, Adam Lerner, and Jeffrey Sanford Russell *

June 2022

Suppose that, for reasons of animal welfare, it would be better if everyone stopped eating chicken. Does it follow that you should stop eating chicken? Proponents of the "inefficacy objection" argue that, due to the scale and complexity of markets, the expected effects of your chicken purchases are negligible. So the expected effects of eating chicken do not make it wrong.

We argue that this objection does not succeed, in two steps. First, empirical data about chicken production tells us that the expected effects of consuming *many* chickens are not negligible. Second, this implies that the expected effect of consuming one chicken is ordinarily not negligible. *Parity* between your purchase and other counterfactual purchases, and *uncertainty* about others' consumption behavior, each tend to pull the expected effect of a single purchase toward the average large scale effect. While some purchases do have negligible expected effects, many do not.

1 Introduction

Chicken consumption is at an all-time high: the average consumer in the United States goes through 94.3 pounds of chicken flesh per year. Although producers have implemented a variety of techniques to increase the amount of flesh that can be harvested per bird, producers have had to raise and slaughter more chickens each year to meet the ever growing demand.

The innovations that have allowed each chicken to yield more meat have come at significant cost to the chickens themselves. As a result, the average

---

*Thanks to Andrew Bacon, Christian Tarsney, the Big Decisions working group at USC, the Pascalian Risks seminar at the Global Priorities Institute, and several anonymous referees for helpful discussion.

broiler chicken's life is dominated by suffering. According to many philosophers, the benefits of eating a chicken's flesh are far smaller than the burdens that the chicken experiences during its short, miserable life. As a result, they think that a world where fewer chickens were raised on factory farms, and people ate vegetarian alternatives instead, would be better.

Nevertheless, many of these philosophers also believe that individuals have no moral obligation to refrain from purchasing chicken flesh. And that isn't because they think eliminating factory-farmed chickens would come at some additional cost beyond reducing people's gustatory enjoyment. It's because they think an individual's chicken consumption is too small to make a difference to the total number of factory-farmed chickens raised and slaughtered. Markets respond to shifts in aggregate demand, not shifts in any individual's demand. If buying an additional chicken cannot increase the number of factory-farmed chickens who are born to suffer and die, then no individual has an obligation to refrain from eating factory-farmed chickens stemming from suffering they cause to chickens. A similar conclusion holds if eating a factory-farmed chicken is sufficiently unlikely to cause any suffering (Budolfson 2019; Nefsky 2011; Harris and Galvin 2012; Fischer 2019).

The choice of whether to buy chicken is one of a ubiquitous class of structurally similar cases. In these, many people face a choice between two options, A and B. It's worse if everyone chooses A than if everyone chooses B. Nonetheless, it seems that any individual choice of A is too unlikely to contribute to the collective bad effects of *everyone* choosing A to generate an obligation on *anyone* to refrain from choosing it. In all such cases, the so-called "inefficacy objection" arises. Here, though, we focus on chickens.[1]

We argue that the expected effect of purchasing a factory-farmed chicken in the marketplace is, as an empirical matter, similar in size to directly causing one factory-farmed chicken to exist and suffer. Thus, if having expected consequences that outweigh its expected benefits makes an action wrong, and if causing a chicken to be born to a life of suffering is indeed much worse than foregoing the benefits that result from eating its flesh, then buying a factory-farmed chicken is wrong.

---

[1] A parallel debate has emerged regarding the individual contributions to climate change. See, for example Sinnott-Armstrong 2005; Kingston and Sinnott-Armstrong 2018; Broome 2019.

2 Thresholds, Expectations, and Inefficacy

2.1 Simple threshold models

Here's a model: every time someone eats a factory-farmed chicken, a signal is sent to to a factory farm to produce another chicken.[2] In this model, the consumption of a chicken obviously causes a chicken to be brought into existence. Assuming that the chicken's lifetime of suffering outweighs your gustatory enjoyment, the consequentialist case against eating a chicken is straightforward.

But this model is obviously unrealistic. A factory farm's output is not set chicken by chicken. Here's another model: factory farms produce and sell chickens in larger blocs. A factory farm might choose between producing either 703,100 chickens or 703,200 chickens in a given year, but it won't choose between producing either 703,148 chickens or 703,149 chickens. If you don't affect whether the factory farm produces either 703,100 chickens or 703,200 chickens then you don't make any difference at all. But if you do affect whether the factory farm produces either 703,100 chickens or 703,200 chickens then you make a big difference.

Let's suppose that our moral reason to avoid a 1% probability of causing 100 chickens to suffer is just as strong as our moral reason to avoid a 100% probability of causing 1 chicken to suffer. Then, if you have a 1% chance of triggering a change from 703,100 chickens to 703,200, the consequentialist case against buying a chicken is still straightforward. Such an appeal to expected consequences has been prominently defended by Peter Singer (1980), Alastair Norcross (2004), and Shelly Kagan (2011).[3] For example, Kagan (2011, 124) reasons as follows:

> [S]ince the butcher neither wants to fall behind demand nor end up with ever larger numbers of unsold rotting chickens, we know as well that the number of chickens he orders is more or less the same as the number of purchases required before a new order is triggered.
>
> Thus we know that there is some triggering number, T (more or less), such that every Tth purchase (more or less) triggers the order of another T chickens (more or less). I don't have any idea

---

[2] For simplicity, we will speak as if people only bought whole chickens; but of course many chickens are instead cut into parts and sold that way, or as ingredients in other foods.

[3] For similar appeals in other contexts, see Gibbard 1990; Parfit 1986; Broome 2019; for further references and discussion see Nefsky 2019.

> what that number is, but I do know that whatever it is, I have a 1 in T chance (more or less) of triggering the suffering of another T chickens (more or less) …
>
> As I walk to the butcher counter, then, not only don't I know whether my act will have bad results, I don't even know what the *chances* are that my act is a triggering act. But I do know, for all that, that the net expected results of my act are bad.

If correct, Kagan's argument would show that the consequentialist case against chicken-consumption goes through whenever there are thresholds of chicken consumption that, when crossed, trigger corresponding increases in chicken production.

But Kagan's point is not right in general. Even if there are thresholds with the structure he assumes, there is no guarantee that that there is some n such that the probability of making a difference to n chickens is $\frac{1}{n}$. If a factory farm scales its production in blocs of n chickens, the probability that eating a single chicken will cause n chickens to suffer can be less than $\frac{1}{n}$ or greater than $\frac{1}{n}$, depending on circumstances. Mark Budolfson (2019, 1716) gives an influential example:

> Richard makes paper T-shirts … 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 far more T-shirts than he is likely to need each month … For many years Richard has always sold between 14,000 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 18,000 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.

In this case, there is a chance that a purchase will trigger production of $5{,}000$ additional T-shirts. But Budolfson correctly argues that, in light of the facts about how Richard operates, a single consumer's chance of triggering such an increase is "dramatically lower than $\frac{1}{5000}$ or any other number that would drive the expected effect of an individual buying 1 T-shirt anywhere near the consequence that 1 additional T-shirt is produced" (2019, 1717).

Chignell (2020, p. 219–220) gives a similar example involving poultry production. He imagines a "Kantian Chicken Factory".

> [T]he regional KCF has a policy of ordering 1,000,000 chickens every month from its supplier in order to meet an average monthly demand just shy of 20 million chicken sandwiches. More specifically, the policy says that if the demand in any given month is between 19,900,000 and 20 million sandwiches, then KCF won't change its usual order … Moreover, average monthly demand in this region has been within these two thresholds every month for the past ten years. Still, the policy is not entirely insensitive to market changes. If the number of sandwich orders in a given month falls below the 19,900,000 threshold, this will trigger a 'lump' reduction in their order for the following month: they will order 995,000 chickens instead of 1 million from the supplier. They have similar thresholds every 100,000 sandwiches below that.

If you buy a chicken sandwich at KCF, then the probability that your purchase will make any difference at all to chicken suffering is substantially smaller than $1/100{,}000$.[4] Moreover, Chignell claims that the world of KCF "is very much like our own."

The mere existence of thresholds designed to keep supply in step with demand does not guarantee the probability of making n units of difference will be anything like $\frac{1}{n}$.

2.2 The inefficacy objection

In general, then, the probability that buying a chicken will make n chickens of difference does not have to be anywhere near $\frac{1}{n}$. This probability is a contingent matter, depending on what kind of informational situation consumers find themselves in.

Proponents of the inefficacy objection contend that facts about modern supply chains give us reason to think that we're never near any threshold that would trigger the production of more chickens in factory farms, and thus that the probability of causing the production of n chickens should be substantially less than $\frac{1}{n}$. Budolfson (2015, 86–87) writes,

---

[4] Even so, the expected effect is larger than you might think: we calculate > 0.08 chickens produced per chicken purchased (that is, per 20 sandwiches). (Details omitted for lack of space.) If the harm to the chicken is > 12 times greater than the benefit of eating twenty chicken sandwiches, then even in this case the sandwiches lose the contest. But we can bolster Chignell's artificial example by imagining an even stronger track record: say KCF's monthly order hasn't changed for a hundred years, rather than ten.

> [M]any 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 arguably 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.

Likewise, Julia Nefsky (2011, 370) writes:

> … Kagan needs that the *factory farm* will increase or decrease future production by *N* chickens depending on whether or not one particular butcher orders another *N* chickens. This seems highly unlikely. 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.

Chickens pass through many hands on their way to consumers' plates. They begin life in hatcheries, are fattened on factory farms, and meet their ends in centralized processing plants. Processors sell the chicken carcasses to wholesalers, who then sell them to retailers, who finally sell them to consumers. According to Budolfson and Nefsky, it is extremely unlikely that information about an individual chicken-purchase would percolate back through this complex system in a way that would impact chicken-production.

These arguments have been endorsed by many philosophers. Some of these philosophers go on to conclude that meat-eating is not morally wrong. For example, Elizabeth Harman (2015, 18) takes Budolfson's argument to constitute a "serious worry" for the view that eating meat is wrong, and builds on Budolfson's reasoning to instead contend that eating meat is a morally permissible moral mistake. But the inefficacy objection has also spawned a cottage industry of philosophers offering alternative arguments that meat-eating is wrong, without relying on the proposition that individual choices have any chance of making a meaningful difference. For example, Julia Nefsky (2017, 2748n13) writes that Budfolson's argument "convincingly demonstrates" that the expected effect of purchasing a chicken is not equal to the average effect of all chicken purchases, and she argues that choosing veganism might even so "help prevent" harm in a different sense. Andrew Chignell claims that Budolfson has "pointed out" that individual meat purchases make no difference, and explores the possibility that eating meat might be wrongful because it involves opportunistically benefiting from others' suffering (2015, 192), or that abstaining might make room for *hope* that fewer animals will suffer, by providing evidence that others might also abstain (2020). Similarly, Tristram McPherson (2015, 88n9) calls Budolfson's argument an "important challenge" to arguments against meat-eating, but argues that we should not eat animal products because it would be cooperating with the wrongful elements of the animal product industry's plans. For similar reasons, Moti Gorin (2017) argues that people who condemn factory farming ought to go vegan to avoid harming *themselves*.

We argue that these philosophers have been too quick to abandon the possibility that eating meat is wrong in virtue of its expected effects on chickens.

3 An Argument for Efficacy

Let us say that purchasing n chickens is *efficacious* when it increases the expected number of factory-farmed chickens by some quantity that is not dramatically less than n. According to the inefficacy objection, buying one chicken is almost never efficacious. We will argue the contrary: indeed, in ordinary circumstances buying a single chicken is often efficacious. The argument for this conclusion rests on two premises.

(1) In ordinary circumstances, buying *many* chickens is efficacious.

(2) In ordinary circumstances, if buying many chickens is efficacious, then buying one chicken is often efficacious.

We will defend each premise in turn—along the way, making each of them more precise.

3.1 Buying many chickens is efficacious

The case that large changes in chicken consumption make large differences to chicken production is empirical, and it is fairly decisive.

In recent decades the production of chicken has been steadily increasing, tracking increases in chicken consumption. In 1997, 8.32 billion broiler-type eggs were hatched in the United States. In 2007, 9.57 billion; in 2017, 9.62 billion. The U.S. population was 19% larger in 2017 than it was in 1997; and the average American ate 28% more chicken in 2017 than in 1997.[5]

When companies decide how many broilers to produce, their decisions are sensitive to whether people are buying huge numbers of additional chickens. For example, we see wholesale warehouses set up their own farms to keep up with demand for their loss-leading rotisserie chickens (Gerlock 2018). Multinational corporations, "scrambling to keep up" with demand, open new chicken processing plants (Durisin 2018). Companies are also ready to scale back if demand should decrease. For example, when restaurant sales suddenly faltered in the midst of the COVID-19 pandemic, the *New York Times* reported: "A single chicken processor is smashing 750,000 unhatched eggs every week" (Yaffe-Bellany and Corkery 2020).

This is unsurprising. If demand for chicken goes up enough, more chickens will be produced. If demand for chicken goes down enough, fewer chickens will be produced. The *size* of these effects is not entirely straightforward, though.[6] If a million new consumers arrive on the scene each purchasing one chicken a week, then ordinarily the price of chicken will rise somewhat; this will lead to a compensating decrease in chicken consumption by other buyers who are price-shopping; and the new equilibrium consumption level will be higher than the original level by something *less* than a million chickens per week. There are other feedback effects as well: for example, a decrease in chicken production will lower the price of chicken feed somewhat, slowing further reductions in production.

Cumulative elasticity is a measure of how much the market equilibrium quantity of a good changes as a result of changes in demand: in our application, this should reasonably approximate the expected number of additional chickens produced per additional chicken consumed. Norwood, Lusk, et al. (2011, 223) report an empirical estimate of the cumulative elasticity of chicken at 0.76. This predicts that if an additional 1 million chickens were consumed, about 760,000 additional chickens would be produced. The

---

[5] These statistics are from the United States Department of Agriculture (USDA) National Agricultural Statistics Service and the USDA Economic Research Service.

[6] Thanks to an anonymous referee for pressing this.

exact value is not essential: the important point is that empirical research strongly supports the view that the expected effect of buying a large number of chickens is not orders of magnitude smaller than one chicken produced per chicken consumed.

To be clear, what we are imagining is not the outlandish counterfactual where a million people all show up at once to buy extra chickens at the same store or even in the same city, emptying supermarket shelves and creating a commotion. Rather, we are imagining what would ensue if there were a million additional chicken purchases scattered throughout the United States at different times through an ordinary week. This would be an increase in total US chicken consumption of about half a percent. We take it that, as an empirical matter, it isn't just dramatic spikes in consumption and local shortages that lead to changes in chicken production with roughly the observed magnitude. Slower, dispersed, unremarkable changes in consumption also lead to increases in production.

Note that the inefficacy objection concerns the workings of actual modern market economies. (We have focused on the United States.) There are other possible economic situations in which things work very differently—and in such circumstances production might not be sensitive to even very large changes in demand. For example, a government might have a policy of buying up excess supply and destroying it in order to prop up food prices, in which case supply might routinely far exceed demand (see McMullen and Halteman 2018, 101). In such circumstances one might buy millions of chickens without risking any effect on how many chickens suffer. But we won't count conditions like these as "ordinary circumstances."

3.2 If buying many chickens is efficacious, then buying one chicken is often efficacious: the parity argument

We will now defend the second premise. We give one argument for it in this section, and a second argument in the next section.

Let the *efficacy* of buying n chickens, E_n, be the expected number of additional chickens that will be raised on factory farms if n additional chickens are bought. We will now take for granted what we argued for in the previous section: specifically, that if a million additional chickens are purchased, in scattered places throughout the United States over the course of a week, then approximately 760,000 additional chickens will be raised on factory farms: that is, $E_{1\text{ million}} \approx 760{,}000$. (Nothing depends on these precise numbers.)

First, then, it must be *possible* for a single purchase to make a difference to expected chicken production. The expected effect of a million purchases is the sum of the expected effect of buying one chicken, the expected effect of then buying a second chicken, the expected effect of buying a third chicken, and so on. (In general, the expectation of a sum of quantities is equal to the sum of the expectations of those quantities.) If the efficacy of buying one million chickens is $760{,}000$, then it can't be that every purchase of a single chicken would have efficacy of less than 0.76. The sum of a million numbers, each of which is less than 0.76, would be less than $760{,}000$. It then follows that if $E_1 \ll 1$, then buying one chicken is much less efficacious than the nth chicken purchase would be, for some n less than a million. (Fact 1 in appendix A makes this more precise.)

In ordinary circumstances, though, it is plausible that a typical purchase you might make is *not* orders of magnitude less efficacious than the second purchase would be, or the third, etc. Call this premise *Parity*. If that's right, then E_1 must not be much less than $\frac{1}{N}E_N$—so it is on the same order as 0.76.

Parity does not *always* hold. An employee at the TigerMarket in Weehauken, New Jersey might know that store's precise ordering policy and monthly sales, and conclude that buying a pack of chicken breasts from that TigerMart would be *especially* unlikely to make any difference to chicken suffering. In particular, the efficacy of her purchase might be much *lower* than the efficacy would be for additionally buying a pack of chicken wings from the Safeway in Mesa, Arizona. But this kind of thing can't be *guaranteed* to happen. While all chicken purchases might be *quite* unlikely to affect chicken suffering, they are not all *especially* unlikely to affect chicken suffering—that is, much less likely than other purchases would be. (Compare Hedden 2020, 538.)

Some defenders of inefficacy will object that Parity hardly *ever* holds in ordinary life.[7] Consider how things work in Budolfson's T-shirt factory or Chignell's KCF (from section 2.1). When circumstances fall in the range of what retailers and producers normally anticipate, an additional purchase makes zero actual difference. Furthermore, in these examples the probability is disproportionately high that circumstances *will* fall in that "normal" range. In that case, if many more chickens *were* purchased than actually *will* be, an additional purchase would be much more likely to make a big difference—so Parity fails.

Take Chignell's chicken factory. Suppose we are very confident that the actual number of chicken sandwiches that will be purchased is close to halfway between 19.9 and 20 million. If one additional sandwich is bought,

---

[7] We thank a referee for forcefully pressing this.

this is very unlikely to make a difference to how many chickens the KCF orders—much less likely than $1/100{,}000$. But the probability is substantially greater than $1/100{,}000$ that, if 50,000 additional sandwiches *were* bought, counterfactually, then an additional chicken purchase would push the total above 20 million. The core of the inefficacy objection is the claim that ordinary consumers are nearly always in this kind of situation.

We maintain that this is *not* in fact the usual situation (at least not to the extent that the inefficacy objection supposes), and that Parity is empirically defensible. Parity says that the efficacy of buying one chicken is about the same as the difference in efficacy between buying $n + 1$ chickens and that of buying n chickens, for each $n < 1$ million. The basic reason this is plausible is because it is very strange to think that buying lots of chickens at other stores around the country, generally quite far away, would matter a lot for the efficacy of this purchase right here. Even though a million chickens in a week would be a large enough shift in aggregate demand to be notable, it would make a very small difference to what is going on at any particular store. It would amount to, on average, fewer than three additional purchases per supermarket in that week. The average supermarket sells chicken products derived from more than a thousand chickens each week. So if the efficacy of one purchase is dramatically different from what it would be if up to a million more chickens were bought across the United States, then either (a) the efficacy would also be dramatically different if *three* more chickens were bought this week at this particular store, or else (b) the efficacy of this purchase is highly sensitive to what is going on at thousands of other far-flung supermarkets. The kinds of consideration that motivate the inefficacy objection do not give us reason to believe either of these striking conclusions.

To simplify things, imagine a world with a million identical retailers, each operating independently, and imagine $n < 1$ million purchases are made, each at a different retailer. Suppose that a purchase of a single chicken at any of these retailers makes a difference of r chickens in expectation. In that world, the efficacy of each of the n purchases is just r—and so exactly the same for each purchase, satisfying Parity. At the other extreme, Parity would be violated if the whole chicken economy was just a scaled up version of one of microcosms that Budolfson, Chignell, and others envision. Neither picture is realistic; but reality's complex markets and supply chains are more like the toy world of a million retailers than they are like a toy world of one.

3.3 If buying many chickens is efficacious, then buying one chicken is often efficacious: the uncertainty argument

The argument of the previous section relied on the Parity premise, that there are not large swings in the efficacy of a single purchase as more chickens are bought, within a reasonable range. We have defended this premise; but it is instructive to consider a second argument for the efficacy of a single purchase which does not rely on any such premise.

Even a very informed consumer should have a fairly wide range of uncertainty about the total level of chicken consumption: we do not have a super precise fix on the state of the agricultural economy. In 2022, we can confidently predict that between 9 and 10 billion chickens will be consumed in the US; but it would be unreasonable for a consumer to be very confident that it will be 9.601 billion rather than 9.602 billion.[8] On scales on the order of a few million chickens, the probabilities of particular total consumption numbers should be approximately uniform.

For illustration, let's start by making the unrealistic simplifying assumption that the probability of n chickens being consumed (other than one you might purchase yourself) is given by a *uniform* distribution between 9 and 10 billion. Buying a single chicken changes the total level of consumption from n to $n + 1$, where n is whatever unknown number of chickens would be consumed if you didn't make the extra purchase. So the *efficacy* of a single purchase can be calculated as the average of various *conditional* efficacies (using the law of total expectation): the expected difference in chicken production supposing that 9 billion and one chickens are consumed, rather than 9 billion; the expected difference in production supposing 9 billion and two, rather than 9 billion and one; and so on all the way up to 10 billion. The average of all of these numbers can be straightforwardly estimated as the expected difference in production if 10 billion chickens are consumed rather than 9 billion, divided by a billion. That is, the efficacy of a single purchase is the efficacy of buying a billion chickens, averaged across all of these purchases.

We emphasize that this argument makes *no* "smoothness" assumptions (like Parity) about *how* expected chicken production depends on the total number of chickens consumed. It doesn't rely on anything like a simple threshold model. Maybe there are large "buffers", reflected by large jumps in expectation at far-apart values of n. We just require that expected chicken production *somehow* ends up at a much larger value when a billion extra

---

[8] McMullen and Halteman (2018) give a related (informal) argument based on uncertainty about production costs.

chickens are consumed.

Instead, this argument uses a "smoothness" premise about reasonable *uncertainty* about total chicken consumption. We don't take this premise to be threatened by the usual considerations that motivate the inefficacy objection. The point is that, even if your expectations about how chicken production depends on chicken consumption are not at all smooth, but rather some very jagged or spiky curve, you don't know where exactly you are on that curve. Your uncertainty about how many chickens will be consumed smooths out the jags and spikes, pulling your expected individual impact closer to the large-scale average.

We illustrated the point using uniform probabilities, but the argument does not depend on this unrealistic assumption: the point holds for any probability distribution that is sufficiently "spread out". This idea is made precise in appendix A (see Fact 2), but we can convey the gist here. We introduce a quantity we call the *N-tightness* of a probability distribution, which typically corresponds to the probability of being within an interval of length N around the distribution's peak. We then show that the difference between the efficacy of a single purchase and the average efficacy of many purchases is constrained by the N-tightness of the probability distribution over total chicken consumption.

We can illustrate this with a simple, though not entirely unrealistic calculation. Suppose our uncertainty about chicken consumption for 2022 can be approximated by a normal distribution with its peak at 9.6 billion chickens, and with a standard deviation of 10 million (figure 1). (This means we are about 95% confident that the total number of chickens consumed will be between 9.58 billion and 9.62 billion. If we are more uncertain than this, the resulting lower bound on E_1 will be even higher, which would strengthen the case against consuming chickens.) For $N = 1$ million, the N-tightness of this distribution is 0.04. If the expected effect of purchasing one million chickens is that $760{,}000$ additional chickens would suffer, then Fact 2 tells us that the efficacy of a single purchase is at least[9]

$$0.76 - 0.04 \cdot 2 = 0.68 \text{ chickens}$$

This estimate should not be taken literally, but the case is strong that the efficacy of purchasing a chicken is not several orders of magnitude smaller than the efficacy per chicken of purchasing many chickens.

Of course, we have acknowledged that there are possible cases, like the T-shirt factory or KCF, where the expected individual effect is *not* all that

---

[9] Fact 2 also depends on a parameter b, which we here set to 2: that is, the expected effect of purchasing one million chickens is at most two million.

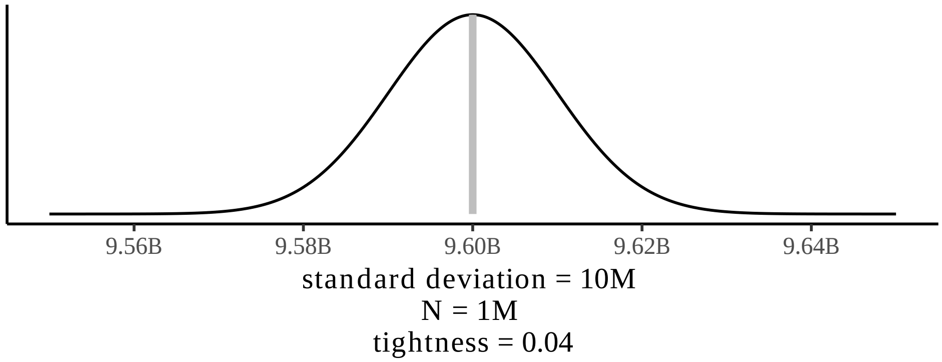

Figure 1: A distribution over chicken consumption

close to the large scale average. How do those cases escape this argument? In the case of KCF, the range of uncertainty about chicken consumption was actually quite small, compared to the *change* in chicken consumption that it would take to make a conspicuous difference in consumption. If we set the probabilities in a way that is favorable to the inefficacy objection, so that individual efficacy comes out < 0.001, this corresponds to being $> 99\%$ confident that between 996,000 and 999,000 chickens will be consumed (in the form of chicken sandwiches). But consumption would have to fall below 995,000 or above 1,000,000 in order to trigger any change in purchasing. In these cases, you *do* have very fine-grained information about others' purchasing behavior, relative to the scale at which there are predictable effects on production. But this is not a realistic model of our actual world. It is implausible that a shift in demand of a million chickens per week (a 0.5% change) is not large enough for economists' methods for estimating cumulative elasticity to be approximately correct. But it is also implausible that a consumer should be very confident of the true level of total consumption to a precision of better than 99%.[10]

---

[10] The uncertainty argument also depends on a subtle independence assumption: namely, that the probability of various dependency hypotheses ("if *n* are consumed, *k* will be produced") are probabilistically independent of how many chickens are actually consumed. This is not necessarily true. For example:

> *Newcomb's Chicken.* An extremely reliable predictor has analyzed your social media history and predicted whether you will purchase a chicken. If they predict you will, they produce two. If they predict you won't, they produce one.

Given the information that you *won't* purchase a chicken, it is very likely that the predictor

4 Objections and Replies

The upshot of these two arguments is that—given that a large number of scattered purchases have proportionately significant effects in expectation—it is not overwhelmingly likely that a single purchase is "buffered" from the rest of the causal chain leading to the factory farm.

This is not to deny Budolfson (2019, 1717) and others' observations about "slack" in this causal chain—for example, "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." We acknowledge that changes in production do not come from some "frictionless optimization procedure" (2015, 88). But it does not follow that this kind of slack creates "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" Budolfson (2019, 1717).

Consider two simple grocery stores. Frugalmart reviews its chicken sales each month, and orders the exact number of chickens they expect to sell the next month. ProdigalCo instead orders the number they expect to sell plus a thousand; chickens unsold at the meat counter are destined for rotisserie, dog food, or the dumpster. The expected effect on chicken welfare of a purchase at ProdigalCo is no better than at Frugalmart. ProdigalCo orders more chickens than Frugalmart overall, but *changes* in their monthly order depend on monthly sales in just the same way. (We are not saying that this simple model of "slack" is realistic; we are using it to show that large wasteful "buffers" don't *have* to decrease the efficacy of a single purchase.)

Budolfson also appeals to the fact that there are many intervening steps between the consumer and "the far other end of the supply chain." He claims that "as long as we can know—as we can—that there are sure to be buffers of non-trivial size throughout the supply chain …that reduces the probability of a single individual making a difference to a level that quickly becomes nearly infinitesimal."[11] Each step in the supply chain has some

---

predicted that you wouldn't buy a chicken, and so produced just one. But also, given this information, it is very likely that the predictor *would have* produced just one chicken even if, improbably, their prediction was wrong and you did purchase one after all. Thus, given this information, the conditional efficacy of your purchase is close to zero. Similar reasoning applies to the case where you *do* in fact make the purchase.

11. Budolfson draws an analogy to voting, claiming that the probability of casting a decisive vote in a large election is "nearly infinitesimal". We are persuaded by Barnett 2020 that this is not true: in competitive elections, that probability is often $> 1/N$, where N is the total number of voters.

significant chance of blocking an effect from further up the supply chain, and these chances compound. Thus, Budolfson claims, "the signal generated by a single individual will almost certainly be lost in transmission and absorbed by buffers" (2019, 1717, 1721).

But if signals really are lost across a complicated chain of transmission in a way that drives down the efficacy of a single purchase, this would also drive down the efficacy of many purchases. Imagine that each retailer sends their order changes to the wholesaler by carrier pigeon, with only a 10% chance of success, and each wholesaler sends *their* order changes to the factory farm by skywriting, which again succeeds only 10% of the time. This would indeed decrease the efficacy of a single purchase. But it would drive down the efficacy of *many* purchases by the same factor. (The probability that *at least one* signal makes it through would, of course, be much higher if there are many purchases. But the expected *proportion* of signals to make it through will be just the same as for a single purchase.) Contrapositively, if many purchases *do* make an important difference to production—as we take observation to bear out—then the probability that an individual signal will make it through cannot be "nearly infinitesimal".

Another worry concerns vagueness (see, for example, Nefsky 2011, sec. X). We have taken it for granted that the expected effect of a chicken purchase, whatever it may be, is a real number that we can reason about using standard mathematics, including classical logic. But plausibly the expected effect of a chicken purchase is not a *precise* matter.

Our view is that vagueness does not undermine our arguments. This is because our favored view is that vagueness allows for *precisifications*.[12] While it is an imprecise matter exactly what number the expectation is, some numbers are *not determinately incorrect*. We can run our argument on any one of these precisifications. Furthermore, whatever is true *on every precisification* of a vague claim is determinately true. So while it may be vague what exactly the expected effect *is*, it can still be determinately true that it is not minuscule.

Of course, there are alternative approaches to vagueness. Our plea to defenders of the inefficacy objection who would rely on vagueness is to present us with an alternative model of how vagueness works in this case, which includes enough detail so that we can calculate the (imprecise) expected effect of a single purchase and that of many. Then we can evaluate that alternative model on its own terms.

---

[12] This terminology comes from supervaluationism, but the idea as we apply it here is common to many leading theories of vagueness.

A final objection is that it just seems *obvious* that if you show up to the supermarket on a single occasion and purchase one whole chicken, then (as an anonymous referee put it) "the folks behind the meat counter would do absolutely nothing differently than if you had never shown up." It seems there must be *something* wrong with arguments that purport to show otherwise.

We appreciate the force of this intuition, but do not accept it.

First, it is important to remember that when we say "the expected effect of buying a chicken is significant", that is not to say that it is reasonable to *expect* that the effect of buying a chicken is significant. Indeed, the actual effects of most chicken purchases probably are insignificant. We agree that we have good reason to think that the signal from an individual's purchase will, as Budolfson says, "almost certainly be lost in transmission"—if "almost certainly" means more than, say, 99.9% likely.

But we don't think that reflection on these facts, and on how difficult it seems for a single purchase to make a difference in light of them, can by itself get us to a quantitative estimate of exactly *how* small the tiny probability of making a difference is—and in particular, whether it is small enough to be negligible for an expected value calculation.

When it comes to *this* question, our intuitions are not trustworthy. Indeed, psychologists have found that when we make intuitive judgments about expected effects, we often implicitly round small probabilities down to zero (see Kunreuther, Novemsky, and Kahneman 2001, among others). The key question is whether the probability of making a difference is *negligibly* small, or merely *very* small; this is not to be settled by direct intuition.

We have argued that this probability is not negligible. But it may be hard to imagine how an individual purchase *could* make a difference to chicken production. For a sketch of an answer to this question, we draw on McMullen and Halteman (2018). First:

> The long and complicated supply chains that connect individual farmers to consumers are actually designed to be more responsive and more reliable than shorter, more local relationships. For example, modern grocery stores have check-out procedures that track the sale of each product and automatically order replacements from the parent companies … [T]here is a wealth of evidence that the vertically-integrated agribusiness oligopolies are the most able to respond to individual changes in consumer demand. (McMullen and Halteman 2018, 104; compare Norcross 2020, 164)

Still, even if the signal from individual purchases sometimes makes its way from the supermarket checkout to the poultry farm, most poultry farms simply produce as much as they can under the constraints of available capital and production costs (as Nefsky 2011, 370 noted). So how could small changes in consumption matter? An answer:

> [M]ost suppliers do not change their behavior in response to a price drop, but rather absorb the losses and keep going. The producers that will respond are those on the margin, with the best outside options or the least competitive position in the market. These marginal producers are the ones that will leave the industry in response to lower prices. In a large industry, moreover, competition will always drive the price to the point where the marginal producer is right on the edge of non-participation in the market. (McMullen and Halteman 2018, 101)

So here is a rough picture of how things might work. Sometimes, if you hadn't purchased a particular chicken at the supermarket, a smaller number of purchases would would have been recorded by a national distributor. Sometimes this would have led to something different being written in a marginal poultry farm's account ledger—perhaps a decrease in the price their processing plant pays per chicken. And sometimes this difference would have made a difference to how long the marginal farm remained operational. Such a chain of events is indeed quite unlikely for any particular purchase—but, as it turns out, not "nearly infinitesimally" unlikely. And one operational day for one poultry farm can mean a day of suffering for 40,000 birds.

5 Conclusion

We have argued that individual consumer choices can make a significant difference to chicken welfare, in expectation, in ordinary circumstances. But, unlike in Singer, Norcross, and Kagan's original consequentialist arguments, our conclusion is *not* that the expected effect of purchasing a chicken is that precisely one additional chicken suffers. These arguments relied on a very simple model, in which the probability that the purchase of a chicken will cause n additional chickens to suffer is $\frac{1}{n}$. The proponents of the inefficacy objection correctly appreciated that this is not realistic. However, we have argued that these proponents were mistaken in concluding that, in real-world cases, the expected effect of buying a chicken is orders of magnitude smaller than the production of one additional chicken. Ordinarily, the expectation is indeed plausibly somewhat less than one, but not negligibly small. (In section 3.3 we estimated > 0.6.)

The exact efficacy value would only make a difference to the consequentialist argument if the trade-off between chicken welfare and the benefits of eating chicken were a pretty close call—which is implausible. If the efficacy of a single purchase was only, say, 0.25, this would only affect the consequentialist calculus in the same way that reducing the time chickens spent in factory farms by a factor of 4 would (given plausible assumptions about welfare aggregation). The typical chicken spends 40 days on a factory farm. If the typical chicken instead spent 10 days on a factory farm, this would not tip the moral balance.

We have focused on choices to purchase chicken, but these lessons generalize. They generalize most immediately to other consumer choices involving a product that the world would be better off without, where large increases in consumption have very bad effects in expectation, and where one does not have detailed information about where one's own contribution stands with respect to the overall pattern of consumption. For example, this may be true of products that contribute to climate change. Saying exactly which kinds of choices have the structure we have described lies beyond the scope of this essay. But for a choice that does have this structure, we can say that its expected consequences are bad.

A Technical Details

In this appendix we sharpen the arguments from parity and uncertainty in sections 3.2 and 3.3.

Let $C(n)$ be the expected number of chickens that will be raised in factory farms if n chickens, in all, are purchased this year. Let $P(n)$ be the probability that n other chickens will be purchased this year (not including any you are currently considering purchasing). Then we can write the efficacy of purchasing k additional chickens as

$$E_k = \sum_{n=0}^{\infty} P(n)\big(C(n+k) - C(n)\big)$$

This formula relies on the simplifying assumption that the k "additional" chickens you are considering purchasing are not special, compared to other chickens that might have been purchased instead: given that n other chickens are purchased, if k additional chickens were purchased, then the expected effect on chicken suffering would be the same as if just $n + k$ chickens had been purchased. Essentially, this amounts to supposing that just the number of chickens that are consumed matters for chicken welfare, not which particular chickens are consumed.[13]

First we restate the argument of section 3.2. Let $d(k)$ be the expected effect of purchasing one further chicken, if $k - 1$ other additional chickens are also purchased:

$$d(k) = E_k - E_{k-1}$$

The Parity premise says that the expected effect of buying this one chicken, $d(1)$, is not much smaller than the expected effect $d(k)$ that buying a kth chicken would have, for $k \leq N$.

Fact 1. *For any number $\alpha > 0$, if $E_1 = d(1) \geq \alpha\, d(k)$ for each $k \leq N$, then $E_1 \geq \alpha \frac{1}{N} E_N$.*

So if the efficacy of buying one chicken would not be dramatically smaller if many other chickens purchased as well, then it is also not dramatically smaller than the efficacy of buying many chickens, averaged across each purchase.

Proof of Fact 1. Since $d(k) = E_k - E_{k-1}$, it follows that

$$E_N = \sum_{k=1}^{N} d(k) \tag{1}$$

(The right-hand side is a "telescoping sum", in which all but the two terms E_N and $E_0 = 0$ cancel out.) If $E_1 \geq \alpha\, d(k)$ for each $k \leq N$, then

$$E_1 \geq \frac{1}{N} \sum_k \alpha\, d(k) = \frac{1}{N} \alpha E_N \qquad \square$$

---

[13] As we noted in footnote 10, it also relies on the simplifying assumption that how many chickens are purchased is probabilistically independent of the dependency hypotheses "if n chickens were purchased, then k chickens would suffer." Given these two assumptions, we can derive the formula for E_k by unpacking these definitions:

$$\begin{aligned} C(n) &= \sum_k \Pr(n \text{ chickens are purchased} \mathrel{\Box\!\!\rightarrow} k \text{ chickens suffer}) \cdot k \\ E_n &= \sum_k \Pr(\text{I purchase } n \text{ chickens} \mathrel{\Box\!\!\rightarrow} k \text{ chickens suffer}) \cdot k \\ &\quad - \sum_k \Pr(\text{I purchase no chickens} \mathrel{\Box\!\!\rightarrow} k \text{ chickens suffer}) \cdot k \end{aligned}$$

The derivation is omitted for reasons of space.

Next, the uncertainty argument. In section 3.3 we made the simplifying assumption that the probability distribution over chicken consumption was *uniform*. This assumption is not realistic, but the idea generalizes. Since we normally don't have extremely fine-grained information about chicken consumption, a realistic probability distribution P will be reasonably "flat", without sharp peaks on a scale smaller than a few million chickens; but a few million chickens is enough to matter.

We can make this precise with a measure of how widely spread out a probability distribution is. Let the *N-tightness* of P be

$$1 - \sum_{n=0}^{\infty} \min_{k<N} P(n+k)$$

This number is close to zero if the probabilities of two numbers within N of each other are always very close together. It is close to one if it has peaks with width less than N. For a unimodal distribution with a single peak, the N-tightness is the same as the total probability that falls in an interval of length N around the peak. So, for example, for $N = 100$, the N-tightness of a normal distribution with a standard deviation of 100 is ≈ 0.34. The N-tightness of a wider normal distribution with a standard deviation of 500 is ≈ 0.08. In general, the distribution P will have a small N-tightness value insofar as you do not have very fine-grained information about the background level of chicken consumption. (See figure 2.)

We can use "tightness" to put another bound on how far the efficacy of a single chicken purchase can depart from the average efficacy of many chicken purchases.

Fact 2. *Suppose that there is an upper bound on the expected effect of consuming N additional chickens: for each n,*

$$C(n+N) - C(n) \leq bN$$

If the N-tightness of P is t, and C is non-decreasing, then

$$E_1 \geq \frac{1}{N} E_N - bt$$

In particular, if P is not too tight and b is not dramatically more than 1, then the efficacy of buying a single chicken, E_1, is close to $\frac{1}{N}$ times the efficacy of buying many chickens, E_N. This result does not depend on any fine details of the function C, which tells us how chicken production depends on chicken consumption.

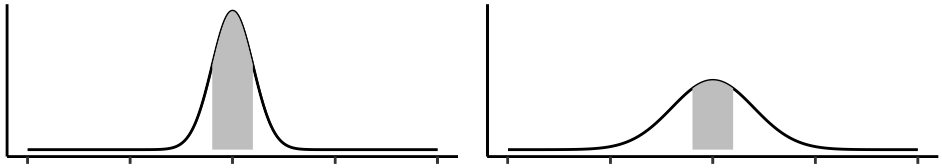

standard deviation = 50
N = 100
tightness = 0.68

standard deviation = 100
N = 100
tightness = 0.38

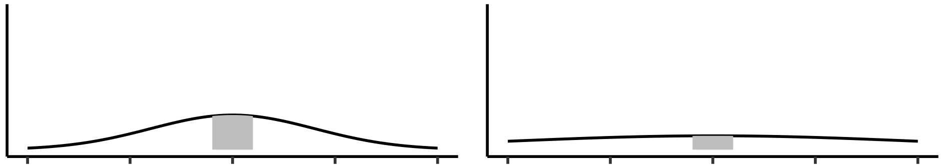

standard deviation = 200
N = 100
tightness = 0.20

standard deviation = 500
N = 100
tightness = 0.08

Figure 2: Tightness values for distributions of different widths

Proof of Fact 2. Let

$$\hat{P}(n) = \min_{k<N} P(n+k)$$

So, by definition,

$$t = 1 - \sum_{n=0}^{\infty} \hat{P}(n) = \sum_{n=0}^{\infty} \big(P(n) - \hat{P}(n)\big)$$

Let $D_k(n) = C(n+k) - C(n)$; so

$$E_k = \sum_{n=0}^{\infty} P(n) D_k(n)$$

For each $k < N$,

$$\begin{aligned} E_1 = \sum_{n=0}^{\infty} P(n) D_1(n) \geq \sum_{n=k}^{\infty} P(n) D_1(n) &= \sum_{n=0}^{\infty} P(n+k) D_1(n+k) \\ &\geq \sum_{n} \hat{P}(n) D_1(n+k) \end{aligned}$$

Adding up these inequalities for each $k < N$, and using the telescoping sum

$$\sum_{k=0}^{N-1} D_1(n+k) = D_N(n)$$

we have

$$\begin{aligned} N \cdot E_1 &\geq \sum_{k=0}^{N-1} \sum_n \hat{P}(n) D_1(n+k) \\ &= \sum_n \hat{P}(n) D_N(n) \\ &= \sum_n P(n) D_N(n) - \sum_n (P(n) - \hat{P}(n)) D_N(n) \end{aligned}$$

Since $D_N(n) \leq bN$, we have

$$\begin{aligned} E_1 &\geq \frac{1}{N} E_N - \sum_n \left(P(n) - \hat{P}(n)\right) b \\ &= \frac{1}{N} E_N - bt \qquad \square \end{aligned}$$

References

Barnett, Zach. 2020. "Why You Should Vote to Change the Outcome." *Philosophy and Public Affairs* 48 (4): 422–446. https://doi.org/10.1111/papa.12177.

Broome, John. 2019. "Against Denialism." *Monist* 102, no. 1 (January): 110–129. https://doi.org/10.1093/monist/ony024.

Budolfson, Mark Bryant. 2015. "Is It Wrong to Eat Meat from Factory Farms? If So, Why?" In *The Moral Complexities of Eating Meat,* edited by Ben Bramble and Bob Fischer. Oxford University Press, November 1, 2015. ISBN: 978-0-19-935393-4. https://doi.org/10.1093/acprof:oso/9780199353903.003.0006. https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199353903.001.0001/acprof-9780199353903-chapter-6.

———. 2019. "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response." *Philosophical Studies* 176, no. 7 (July 1, 2019): 1711–1724. ISSN: 1573-0883. https://doi.org/10.1007/s11098-018-1087-6. https://doi.org/10.1007/s11098-018-1087-6.

Chignell, Andrew. 2015. "Can We Really Vote with our Forks? Opportunism and the Threshold Chicken." In *Philosophy Comes to Dinner: Arguments about the Ethics of Eating,* edited by Andrew Chignell, Terence Cuneo, and Matthew C. Halteman. October 8, 2015. https://doi.org/10.4324/9780203154410-20.

———. 2020. "Hope and Despair at the Kantian Chicken Factory: Moral Arguments About Making a Difference." In *Kant and Animals,* edited by John J. Callanan and Lucy Allais, 213–238. Oxford University Press. https://doi.org/10.1093/oso/9780198859918.003.0012.

Durisin, Megan. 2018. "Can Americans Eat More Chicken? U.S. Industry Bets on Growth." *Bloomberg.*

Fischer, Bob. 2019. *The Ethics of Eating Animals: Usually Bad, Sometimes Wrong, Often Permissible.* Routledge. https://doi.org/10.4324/9780429277887.

Gerlock, Grant. 2018. "Costco Builds Nebraska Supply Chain For Its $5 Rotisserie Chickens." *NPR: All Things Considered.*

Gibbard, Allan. 1990. *Utilitarianism and Coordination.* Garland.

Gorin, Moti. 2017. "Causal Inefficacy and Utilitarian Arguments Against the Consumption of Factory-Farmed Products." *Journal of Agricultural and Environmental Ethics* 30, no. 4 (August 1, 2017): 585–594. ISSN: 1573-322X. https://doi.org/10.1007/s10806-017-9690-7. https://doi.org/10.1007/s10806-017-9690-7.

Harman, Elizabeth. 2015. "Eating Meat as a Morally Permissible Moral Mistake." In *Philosophy Comes to Dinner: Arguments about the Ethics of Eating,* edited by Andrew Chignell, Terence Cuneo, and Matthew C. Halteman. October 8, 2015. https://doi.org/10.4324/9780203154410-20.

Harris, John Richard, and Richard Galvin. 2012. "'Pass the Cocoamone, Please': Causal Impotence, Opportunistic Vegetarianism and Act-Utilitarianism." *Ethics, Policy & Environment* 15, no. 3 (October): 368–383. https://doi.org/10.1080/21550085.2012.730258.

Hedden, Brian. 2020. "Consequentialism and Collective Action." *Ethics* 130, no. 4 (July): 530–554. https://doi.org/10.1086/708535.

Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy & Public Affairs* 39 (2): 105–141. ISSN: 1088-4963. https://doi.org/10.1111/j.1088-4963.2011.01203.x. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1088-4963.2011.01203.x.

Kingston, Ewan, and Walter Sinnott-Armstrong. 2018. “What’s Wrong with Joyguzzling?” *Ethical Theory and Moral Practice* 21 (1): 169–186. https://doi.org/10.1007/s10677-017-9859-1.

Kunreuther, Howard, Nathan Novemsky, and Daniel Kahneman. 2001. “Making Low Probabilities Useful.” *Journal of Risk and Uncertainty* 23, no. 2 (September): 103–120. https://doi.org/10.1023/A:1011111601406.

McMullen, Steven, and Matthew C. Halteman. 2018. “Against Inefficacy Objections: The Real Economic Impact of Individual Consumer Choices on Animal Agriculture.” *Food Ethics* 1 (4): online first. https://doi.org/10.1007/s41055-018-00030-4.

McPherson, Tristram. 2015. “Why I Am a Vegan (and You Should Be One Too).” In *Philosophy Comes to Dinner: Arguments about the Ethics of Eating,* edited by Andrew Chignell, Terence Cuneo, and Matthew C. Halteman. October 8, 2015. https://doi.org/10.4324/9780203154410-20.

Nefsky, Julia. 2011. “Consequentialism and the Problem of Collective Harm: A Reply to Kagan.” *Philosophy & Public Affairs* 39 (4): 364–395. https://doi.org/10.1111/j.1088-4963.2012.01209.x. https://www.jstor.org/stable/23261251.

———. 2017. “How You Can Help, Without Making a Difference.” *Philosophical Studies* 174, no. 11 (November 1, 2017): 2743–2767. https://doi.org/10.1007/s11098-016-0808-y. https://doi.org/10.1007/s11098-016-0808-y.

———. 2019. “Collective Harm and the Inefficacy Problem.” *Philosophy Compass* 14 (4): e12587. ISSN: 1747-9991. https://doi.org/10.1111/phc3.12587. https://onlinelibrary.wiley.com/doi/abs/10.1111/phc3.12587.

Norcross, Alastair. 2004. “Puppies, Pigs, and People: Eating Meat and Marginal Cases.” *Philosophical Perspectives* 18 (1): 229–245. https://doi.org/10.1111/j.1520-8583.2004.00027.x.

———. 2020. “The Impotence of the Causal Impotence Objection.” *Southwest Philosophy Review* 36 (1): 161–168. https://doi.org/10.5840/swphilreview202036118.

Norwood, F. Bailey, Jayson L. Lusk, et al. 2011. *Compassion, by the Pound: the Economics of Farm Animal Welfare.* Oxford University Press. https://doi.org/10.1093/acprof:osobl/9780199551163.001.0001.

Parfit, Derek. 1986. *Reasons and Persons.* Oxford University Press. https://doi.org/10.1093/019824908X.001.0001.

Singer, Peter. 1980. “Utilitarianism and Vegetarianism.” *Philosophy & Public Affairs* 9 (4): 325–337.

Sinnott-Armstrong, Walter. 2005. “It’s Not My Fault: Global Warming and Individual Moral Obligations.” In *Perspectives on Climate Change,* edited by Walter Sinnott-Armstrong and Richard Howarth, 221–253. Elsevier.

Yaffe-Bellany, David, and Michael Corkery. 2020. “Dumped Milk, Smashed Eggs, Plowed Vegetables: Food Waste of the Pandemic.” *New York Times* (April 11, 2020). https://doi.org/10.1016/S1569-3740(05)05013-3.