

Forthcoming in *The Oxford Handbook of Food Ethics*, ed. Anne Barnhill, Mark Budolfson, and Tyler Doggett.

Consumer Choice and Collective Impact¹

Julia Nefsky

1. Introduction

It's 10pm and you are channel surfing. You settle on *Beans of Darkness*, a documentary about injustice in the coffee industry. You learn about grueling working conditions; exploitative wage systems that keep farmers living in poverty and debt; and the use of child and slave labor. The movie ends with a message about our power as consumers to change the situation: by demanding fair trade coffee – it says – we can put an end to these injustices. On your way to work the next morning, you stop at your usual coffee shop. Normally you order the Broadway Blend, but today you study the options. There is one coffee labeled 'fair trade'. You order it. It's decent, but not nearly as good as the Broadway. The following morning, back at the coffee shop, you find yourself unable to motivate ordering the fair trade coffee again. It's not that you no longer care about the lives of coffee farmers. It's that it now seems naïve to think that buying a cup of fair trade coffee could do anything to help them. You are a mere speck in a sea of coffee drinkers. Realistically, nothing in the industry is going to go differently depending on which cup of coffee you buy. What, then, is the point in forgoing the tastier coffee?

Many of us encounter similar barriers when we direct ourselves to make ethically conscious food choices. Taken collectively, our food choices are the driving factor behind what

¹ Many thanks to Mark Budolfson and Tyler Doggett for detailed comments on an earlier draft of this chapter. For helpful comments and discussion, I would also like to thank Jonathan Anomaly, Andrew Franklin-Hall, Chris Melenovsky, Hamish Russell, Walter Sinnott-Armstrong, Olivia Sultanescu, Wayne Sumner, Sergio Tenenbaum, and audiences at the University of North Carolina Chappell Hill's Political Theory Workshop, York University, and the Centre for Ethics at the University of Toronto.

players in the industry do to compete and gain our business. Thus, collectively, they have major consequences for human lives, animal lives and the environment. But it is far from clear what this fact about our collective food choices tells us about any individual choice. When you learn about the suffering inflicted on animals on factory farms, or about the exploitation of farm workers throughout the world, or about devastating effects of industrial farming or food processing on the environment, what should you conclude about what you ought, or ought not, to eat for dinner tonight?

One major problem in answering this question is that it is often difficult to determine which large-scale consumer patterns would result in changes for the better. Determining what the actual effects would be of various large shifts in consumer demand is a complex matter.² But even supposing that we do know which large-scale consumer patterns would be better, a problem arises in trying to move from these facts to conclusions about what one ought, individually, to do. Suppose we know that grave harms or injustices could be prevented if enough people stop buying *Xs*. Intuitively, in such a scenario, morality should yield that each of us ought or, at least, has good moral reason to refrain from buying *Xs*. But the trouble is: typically, in such scenarios, it is doubtful that it will make a difference whether or not *I* purchase an *X* on any given occasion. Most present-day food companies operate at too large a scale for a single purchase to have an effect on what they do. And if this is right, then for each of us, it is unclear what point there is in refraining. Each individual can argue, “my purchase won’t make a difference, so I’m not doing

² There is disagreement, for instance, about the effectiveness of the Fair Trade certification program at reducing poverty among coffee farmers. In this paper I will assume, for the sake of discussion, that a widespread shift in demand to fair trade coffee would be overall better for farm workers, but I note here that this may not be true.

anything wrong.” It is this second problem that I am concerned with in this paper. Call it, “the problem of collective impact.”³

Many of our food choices face this problem in one way or another. For another example, take meat consumption. People often move immediately from the fact that tremendous suffering is inflicted on animals on factory farms, or that factory farming causes serious environmental damage, to the conclusion that it is wrong to purchase meat. But this ignores a major complication: it’s highly doubtful that my purchasing some factory-farmed meat will make a difference with respect to animal suffering or environmental damage.⁴ For my purchase to make a difference with respect to these things, it must make a difference to what happens on the factory farm. And while the factory farm may need to adjust production if there is a large shift in demand, a single purchase more or less is surely not going to affect what they do. For example, in the 2013 fiscal year Tyson produced an average of 40.9 million chickens *per week*.⁵ They are not going to adjust production in virtue of the fact that one, or two, or twenty fewer chickens were purchased than usual.

It is important to see that the “my purchase won’t make a difference” claim challenges not only the thought that I am obligated all-things-considered to refrain from consuming meat but also, and more deeply, the thought that there is any moral point at all in refraining (or, more precisely, any moral point having to do with animal suffering, or the environment.) If my purchase won’t make a difference with respect to animal suffering, this means that things will be exactly the same for animals give or take my purchase. And if that is true, it’s not clear how the

³ In previous work I have called this “the problem of collective harm”, but this new terminology better captures the nature of the problem. I am grateful to Andrew Franklin-Hall for this suggestion.

⁴ There is a third issue with factory farming: its implications for public health. Factory farming spreads viruses and antibiotic-resistant bacteria. (Anomaly 2015) The problem of collective impact arises with respect to this issue too.

⁵ Tyson Foods, Inc. Fiscal 2013 Fact Book, http://s1.q4cdn.com/900108309/files/doc_downloads/Tyson%202013%20Fact%20Book.pdf Accessed on 12/24/2015.

suffering of animals on factory farms could give me *any* reason, or presumptive moral obligation, not to purchase it, let alone an all-things-considered obligation.⁶

The problem of collective impact is not unique to food choices. It arises in any context in which by acting in a certain way, we collectively can have a major impact, but in which no individual seems to make a difference.⁷ But, my focus here is on the problem as it arises in the context of food choices. Whether one is coming at this from a general interest in the problem of collective impact or from an interest in the ethics of food choices, it is important to think through how the problem plays out in this domain.

In this paper, I discuss a number of responses to the problem. I do not attempt to cover every potentially promising response. Rather, I focus on a selection of responses that seem to me to be particularly salient in food choices cases. In section 2, I consider two attempts to solve the problem by denying the “it won’t make a difference” claim. In section 3, I discuss appeals to what I call “non-instrumental” moral reasons: reasons that do not have to do with making a significant instrumental contribution. And there I explore a phenomenon that suggests that many people think of the morality of food choices in a not-purely-instrumental way: many people who are vegetarian (for moral reasons) will not eat meat even when they know it would otherwise simply go to waste. I explore whether thinking about the motivations of such vegetarians might

⁶ While I am going to focus on individual purchases, the same problem arises if we ask whether one should adopt a personal food policy. For example, it’s doubtful that your becoming vegetarian will make a difference to the suffering of animals. Suppose you would buy four chickens each month, were you not vegetarian. If you live for sixty more years from the time you become vegetarian, this means 2,880 fewer chickens are purchased in virtue of your being vegetarian. But you can’t figure out whether your being a vegetarian makes a difference by asking, ‘would the factory farm change production if 2,880 fewer chickens were purchased?’ As Shafer-Landau says, “The fact that one has spent ten thousand dollars over a lifetime does not entail that one has exercised during one’s lifetime the causal influence attainable by a one-time ten thousand dollar purchase.” (Shafer-Landau 1994, p. 89) Perhaps the factory farm looks at data about sales on a monthly basis, and decides whether to alter production. A difference in sales of four purchases in a month surely won’t affect their production decisions. But if that’s so, then your being vegetarian does not make a difference to what they do.

⁷ A classic example is voting in a large national election: if enough of us vote for the better candidate, we together can bring it about that she is elected; but your individual vote, you can be sure, is not going to make a difference.

help solve our problem. In the concluding section of the paper, I briefly explain my own view about where the solution to the problem lies.

Before proceeding, a clarification: not just any moral reason or requirement is a candidate solution. Suppose you argue that there is something intrinsically wrong with eating animals, separate from anything to do with animal suffering or environmental damage. This would be a story about why one ought not to eat meat, but it is not even a candidate to solve the problem of collective impact that arises with respect to meat consumption. The problem of collective impact is a problem specifically about the fact that our food choices, taken collectively, can have a major influence on certain morally important outcomes (e.g. animal suffering, worker exploitation, environmental damage.) It seems that this should mean that each of us has strong moral reason to make certain sorts of food choices rather than others, but it's not clear how this can work given the apparent instrumental insignificance of any individual choice. We cannot address this problem by invoking moral considerations that have nothing to do with those outcomes. That would be to just change the topic. I will not always be explicit about this. I will often use "moral reason" as short for 'moral reason that has to do with the outcome of concern'.

2. You *Can* Make a Difference

One common reaction to the "my purchase won't make a difference" claim is to try to deny its truth. If it is not true that it won't make a difference then, presumably, the problem is dissolved. In this section, I consider what seem to be the two most promising ways to take this approach.

2.1 Making a Difference Indirectly

If one is trying to show that a single purchase actually can make a difference, a natural place to look is the influence one's own actions can have on other people. If my purchase

influences others to buy as I do, then whether or not I make a difference does not depend only on whether my choice directly effects a relevant change in production levels or business practices. I will have made a difference if the combination of my choice and the choices of those I influence, taken together, results in a change. Once we take this into account – the claim goes – we see that I very well could make a difference.

There is definitely an important point in this response: we should not forget about the ways our decisions can influence others. This can affect what one should do. However, this will not work as a solution to the problem. The most obvious issue is that it does not apply generally enough. First, for most of us, our individual everyday purchases do not usually influence others to act similarly. More often than not, what I buy in the grocery store is not noticed by anyone, let alone causing anyone to make similar choices. Second, even when a purchase does influence others, unless one is a particularly ‘watched’ person (a celebrity, politician, etc.), it probably does not influence enough others for this influence to provide grounds for thinking that the act made a difference. If my mom and brother are influenced by my decision to eat lentils rather than chicken at the restaurant tonight, and they do the same, it seems just as doubtful that this will make a difference to production on factory farms as the case in which they are not influenced.

There is at least one way in which an ‘ordinary’ person’s purchase could make a difference via its influence. Occasionally, not only does one person influence others, but those people influence further others, and so on. If your influence compounds in this way, it is plausible that the outcome would have been relevantly different had you not acted as you did.⁸ We should take this possibility into account. But we also should not overestimate our ability to

⁸ Norcross makes this point (2004, p. 233), but his main argument is the one I will discuss in 2.2. Glover discusses a slightly different version of compounding influence (1975, pp. 179-181).

have this sort of compounding influence. For most of us, most of our purchases do not set off any such chain-reactions.

Now, what if one says: “I see that you were hoping that there would be a solution that applies to everyone, at all times. But there isn’t. The best we have is that morality tells us to make the choices in question when one could have a substantial influence.” This reply means accepting that many of us have no moral reason, much of the time, to make the choices in question, and that is the main reason to be dissatisfied with it. But it’s also worth noticing what this entails for the *content* of the indirect difference-making reason itself. Suppose that Beyoncé is spotted purchasing a cup of fair trade coffee, and this leads thousands of people to start choosing fair trade. Their taking it that Beyoncé thinks buying fair trade is morally important makes thousands of fans think so too, and thus follow suit. If indirect difference-making were the only answer to the problem of collective impact, then – unlike Beyoncé – most of these fans have no moral reason to choose the fair trade coffee, since most do not have the relevant kind of influence. But this would mean that what the proposal says is that the moral reason Beyoncé had to purchase fair trade was, essentially, that by doing so she could *fool* people into doing something that morality does not recommend that they do – something morally pointless. This is highly unintuitive. It is far from what most of us have in mind when we think of influential people as having moral reason to make certain choices because of their power of influence. Intuitively, what someone like Beyoncé has moral reason to do is use her influence to get others acting in morally good ways – not to fool them into doing things that are pointless. This intuitive conception implicitly relies on there being a separate solution to the problem of collective impact.

We already have enough reason to be unsatisfied with the indirect difference-making proposal. But there is a more basic issue with it – one that sheds light on the search for a solution more generally. To bring this out: imagine that we live in a world in which most people do have the influencing potential necessary for their everyday purchases to frequently make an indirect difference. Even in such a world, I don't think we should be satisfied with the indirect difference-making proposal. In such a world, our influencing potential would give each of us moral reason, or a presumptive moral obligation, not to eat meat, to buy ethically-produced coffee, etc. But it would only do so because of an empirical contingency that is not necessitated or made true by the conditions that give rise to the problem. Such a consideration does not address what is at issue in the problem.

The problem of collective impact is that it seems that

(A) The fact that if enough of us make certain readily available food choices rather than others, grave harms or injustices would be prevented or reduced,

should be enough empirically for morality to yield that

(B) Each of us has a (presumptive) moral obligation, or at least good moral reason, to make those choices.

But it's unclear how this can work, given that

(C) What one does individually doesn't seem to make a difference.

A moral reason that depends on empirical contingencies besides either (A) or the conditions that give rise to (A) does not address that problem. Or, to put it another way: if such a reason *is* your answer, then you are really biting the bullet on the issue. You are accepting that *despite* (A), it could very well be that each of us has no moral reason to make the relevant choices. You are accepting, in other words, that unless certain additional things happen to be true, there would be no moral reason for individuals to refrain from eating factory-farmed meat or to buy fair trade

coffee, etc., *even though* if enough of us do so we could prevent tremendous suffering or injustice. I am not ruling out that biting that bullet could be the right way to go. But we certainly should thoroughly explore non-bullet-biting solutions before we accept such a conclusion.

2.2 Expected Utility

A popular approach to the problem among philosophers is to appeal to chances of making a difference. Those who take this approach argue that in the cases in question, there is always at least *some* chance of making a difference. Now, merely showing that there is some chance of making a difference is not enough. If my eating factory-farmed meat has only a miniscule chance of making a difference for the worse, it is still going to be unclear why morality would require or recommend refraining. After all, almost any action you could perform has at least a miniscule chance of going awry and causing some negative outcome. But what these philosophers argue is that given what we can know about the chances of making a difference, we can see that a calculation of expected utility will yield the moral verdicts we are looking for. Shelly Kagan, Alastair Norcross, and Peter Singer each give arguments of this sort in the case of purchasing factory-farmed chicken.⁹ Their specific presentations vary, but their core argument is the same. Here it is:

Suppose the executives of a factory farm review their numbers – sales, profits, etc.– on a monthly basis. While the company isn't going to adjust production depending on whether one or two fewer chickens were sold this month than projected, there must be some point at which a shortfall in sales would lead them to decrease production. Let's suppose that quantity is 1000 chickens. 1000 chickens is the 'unit' that they make decisions in terms of. For every shortfall of 1000 chickens in a given month, the factory farm will respond by cutting production by 1000

⁹ See Singer 1980, Norcross 2004, and Kagan 2011.

chickens the next month. In this case, each consumer has a $1/1000$ chance of saving 1000 chickens from a life of suffering on the factory farm.

Why a $1/1000$ chance? Because this – the claim goes – is the chance that the shortfall in a given month is an exact multiple of 1000, and one's refraining from buying a chicken makes a difference *if and only if* the shortfall is an exact multiple of 1000. To illustrate the latter claim: suppose the shortfall is exactly 2000, and thus that there is a decrease in production the next month by 2000 chickens. In this case, *had* you purchased a chicken, the shortfall would have been 1999, and the decline in production that would have resulted would have only been 1000. So, your refraining made the difference between 1000 more chickens suffering and not. In contrast, if the shortfall is anything other than a multiple of 1000 then your refraining made no difference at all to chicken production.

The final step of the argument is to show that this $1/1000$ chance, while small, is enough for the expected utility to come out against purchasing a chicken. This is straightforward. A $1/1000$ chance of preventing 1000 chickens from lives of suffering is an expected utility equivalent to a 100% chance of saving a single chicken from that life. By the highly plausible assumption that the benefits obtained by my purchasing a chicken are less substantial than the harms caused to a chicken on a factory farm, it follows that the expected utility of refraining is higher than the expected utility of purchasing a chicken, and thus that – on the expected utility model – I ought to refrain.

As those who advance this argument emphasize, the particular number chosen (in my case, 1000) does not matter. What the argument relies on, though, are the following two claims:

(1) There is some number N , such that for each shortfall in sales of N chickens, the company will decrease production by a corresponding N chickens.¹⁰

(2) The probability that the shortfall will be an exact multiple of N is $1/N$.

The chief problem with the argument is that these claims need not be true, or even close to true.

Regarding (1): there is no reason to assume that the decrease in production triggered by a shortfall will match the size of the shortfall.¹¹ The intuition that they must match comes from the thought that the factory farm will want supply to match demand. But there are many strategies a factory farm could take to bring supply in alignment with demand, including marketing campaigns, searching for new buyers within the market, finding new markets to sell to, and so on. So, even if they take a shortfall of N or higher to indicate that demand has dropped far enough to warrant a decrease in production, they need not take it to warrant or require a decrease by a corresponding N or more chickens. They could, instead, take a “mixed” strategy – e.g. some decrease in production (much less than N) and some expansion into new markets. Since this could very well be how things go, there is no *a priori* reason to think that the expected utility model will yield that you ought not to purchase a chicken. It could easily come out the opposite.

Things are much worse for the expected utility argument than that. Mark Budolfson has persuasively argued that claim (2) not only need not be true, but is in fact false in the world we

¹⁰ It's not strictly speaking true that their argument relies on (1). What it relies on is that there are thresholds at which a decline in sales will trigger a decrease in production, and that the decrease in production triggered will match (or roughly match) the size of the threshold decline in sales. The numbers do not have to be as neat as their actual arguments suggest, and so do not need to be as neat as in (1). But just as their arguments do not strictly-speaking depend on the numbers being as neat as they have them, my objection here does not depend on it. So, there is no loss of generality in putting my critique in terms of (1).

¹¹ I make this point in Nefsky 2011, pp. 369-371. You might think there is a more basic issue: are there really precise numbers at which the executives would make different production decisions give or take a single purchase? I argue in Nefsky 2011 against the claim that in all collective impact cases there must be sharp thresholds like this. But it would not be worth pressing that line here, because I think that most (though, not necessarily all) food choice cases do contain sharp thresholds (though they will typically be unknown, and might vary from moment to moment).

live in.¹² His argument is roughly as follows. First, we know that massive companies like factory farms, and the businesses they sell to, have large ‘buffers’ built in – practices and pricing strategies that ensure that the farm will not have to adjust production based on small-ish deviations in sales. What this means is that the size that a shortfall in sales would have to reach (‘ N ’) for the response to be a decrease in production is going to be *large*. For example, perhaps for every 5% shortfall in overall sales, the factory farm will decrease production.¹³ For Tyson in 2013, this would be an N equal to approximately 8 million chickens. But once we know that we’re talking about numbers like this, we know that the chance of making a difference is not going to be anywhere close to $1/N$. $1/N$ is the chance of making a difference *only if* we assume that any size shortfall is equally likely as any other. But not all consumer behavior is equally likely; it is much more likely that consumers will behave close to how the company projects. And for it to be as much as 8 million purchases off is extremely unlikely. What this means is that the probability of being in the vicinity of a threshold at which they would decrease production is very low, and the probability of *hitting* one of these thresholds exactly, such that a single purchase makes a difference, is “nearly infinitesimal”.¹⁴ It is going to be, in other words, nowhere near what it would need to be for the expected utility model to yield that one ought not to purchase a chicken.

¹² Budolfson Forthcoming.

¹³ While writing this paper, I had the opportunity to speak to an executive at a large agricultural company (who prefers to remain anonymous), and my conversation with him suggested that, at least for some companies, 5% isn’t even high enough. I asked him about whether the current gluten-free diet trend resulted in a decrease in wheat production in his company, and he said “no”. He explained that this trend only resulted in a “5% hit in the market”, and that that small of a decline “doesn’t really matter to them”.

¹⁴ Budolfson Forthcoming, p. 11.

2.3 Conclusion So Far

It is doubtful that we can solve the problem of collective impact in food choice cases by rejecting the “it won’t make a difference” claim: in many of these cases, for many people, one’s purchases effectively won’t make a difference.

A natural place to turn at this point is to what I call “non-instrumental” moral reasons: moral considerations that do not have to do with playing a significant instrumental role in preventing the relevant harms or injustices, and instead concern other ways one’s choice can relate to those outcomes. In section 3 I discuss this strategy, focusing in on two such proposals.

3. Non-Instrumental Reasons

Consider two kinds of vegetarians: those who would eat already-purchased meat if they knew that it would otherwise simply go to waste, and those who would not. If we are looking for non-instrumental moral reasons to refrain from consuming meat, it seems worthwhile to think about the second kind of vegetarian. A vegetarian of the first kind might just be mistakenly assuming that her refusing to eat meat in ‘normal’ cases will actually *do something* to prevent animal suffering. The second kind of vegetarian, on the other hand, clearly has other sorts of reasons in mind.

To make this concrete, take the following case:

Leftovers: You are housesitting for your Aunt Neta, who left town this morning. In her fridge you find some very tasty-looking leftover chicken. You know that these are leftovers from her dinner last night, and thus that they would be safe to eat. But you know that were this food to wait until she came back, it would be spoiled. You are the only one who could eat it. It will merely go to waste if you don’t eat it.

Many, though not all, vegetarians would not eat the leftovers in this case. Some would not eat them because of “non-moral” reasons: for instance, they think it would be unhealthy, or they have no desire to eat it. But I take it that there are vegetarians who would not eat it for moral

reasons, and in particular, reasons having to do with the suffering animals undergo on factory farms. If there are such moral considerations speaking against eating meat in Leftovers, it seems these should be able to provide a relevant answer as to why we should not consume meat in the normal case.

Here is one explanation, though, for not eating the meat in Leftovers that doesn't fit that bill: "if I eat meat in this context, this might lead me to be tempted to do so at other times, when it won't merely go to waste." This story for why I ought not to eat meat in Leftovers only works if (a) it is really true that doing so might lead me to eat meat at other times, and (b) it is really true that I ought not to eat it at those other times. Whether (a) is true will vary individual-by-individual. Whether (b) is true depends on whether there is a *solution* to the problem of collective impact. This explanation, thus, presupposes rather than provides an answer to the problem.

But perhaps other moral reasons for not eating meat in Leftovers will do better.

3.1 Benefiting from Wrongdoing

One explanation is that eating the meat would amount to benefiting from someone else's wrongdoing, and this is something one ought not do. Factory farms wrongfully torture the animals they raise for slaughter. If Aunt Neta's chicken was produced on a factory farm, then if you eat it, you would be voluntarily enjoying a product of this morally abhorrent behaviour. This is something you ought not do. If this can explain why you ought not eat the meat in Leftovers, it will work in normal cases of consuming factory-farmed meat too.¹⁵

¹⁵ Lawford-Smith gives this proposal (2015, pp. 317-318), but it is not her main proposal. I discuss her main proposal below in note 21. Note that the 'benefiting from wrongdoing' proposal is an example of a candidate solution that is available in food-choice cases, but not in various others in which the problem of collective impact arises. It won't work, for instance, to explain why one should vote in the national election: you are not benefiting from someone else's wrongdoing by not voting.

This proposal has some intuitive appeal, but I don't think it works. This is because I don't think it's true that the mere fact that one is benefiting from wrongdoing makes one's action wrong.¹⁶ Begin with this case from Barry and Wiens:

Terrorist Bombing: "A terrorist sets off a bomb, which grievously injures several people. To avoid the explosion, Bill retreats to a nearby café, where he meets Susan. Their chance meeting eventually gives rise to an extremely lucrative business partnership."¹⁷

In engaging in this partnership, Bill and Susan are benefiting from wrongdoing: they wouldn't have met, and thus could not have started this business, had the bombing not occurred. But clearly, in doing so, Bill and Susan are doing nothing wrong.¹⁸ Or, at least, the fact that they would be benefiting from someone else's wrongdoing does nothing to make their engaging in it wrong. (It could be wrong for other reasons). This, at first pass, shows that benefiting from wrongdoing is not in itself wrongful.

Defenders of the proposal might reply, though, that this is not a counterexample because Bill and Susan are *not* benefiting from wrongdoing in the relevant sense. Bill and Susan's benefiting is only an incidental consequence of the wrongdoing: it is mere coincidence that the bombing leads to their meeting. To count as "benefiting from wrongdoing", in the relevant sense – they could say – there must be a non-incidental connection between the wrongdoing and the benefiting.¹⁹

¹⁶ A different worry is that it is not always true that by buying or consuming a product you are *benefiting*. Imagine that you buy some meat, but then throw it out because you end up finding it repulsive. You are clearly not getting any benefit from the purchase. (Thanks to Tyler Doggett for this example.) Perhaps advocates of the proposal could reply that what is wrong is *trying* to benefit while knowing that this benefit would be a result of wrongdoing. Alternatively, they could spell out a technical notion of 'benefiting' such that any purchase or use of something counts as 'benefiting' in the relevant sense. In fact, it seems to me that for this proposed solution to avoid the problem of dependence on additional, contingent empirical facts, they need to take the latter approach.

¹⁷ Barry and Wiens Forthcoming, p. 3.

¹⁸ Barry and Wiens use the case to show that a closely related but distinct view is mistaken: namely, the view that those who benefit from wrongdoing owe special duties of compensation to the victims of that wrongdoing.

¹⁹ Thanks to Mark Budolfson, Tyler Doggett, and Sergio Tenenbaum for raising this point.

But even on this narrower conception, I don't think it's true that the mere fact that one is benefiting from wrongdoing is enough to make one's act wrong. Consider:

Medical Treatment: Beth receives a medical treatment, which saves her life. The discovery of this treatment is owed ultimately to a discovery made by Nazi experiments on concentration camp victims.

Here the connection between the wrongdoing and the benefiting is non-incidental: the Nazi experiments were aimed specifically at advancement in medical treatments. But even if she knew the history of the treatment, it seems clear to me that Beth did nothing wrong in accepting it, and thus in (knowingly) benefiting from wrongdoing. Some, however, might have conflicting intuitions: they might think it is at least *pro tanto* wrong for Beth to accept the treatment. But even if it were *pro tanto* wrong, would it be just the fact that Beth is benefiting from the Nazi's wrongdoing that explains this? There are other more plausible explanations. For instance: perhaps if people like Beth are willing to use this treatment, this could set a risky precedent, potentially encouraging current researchers to violate codes of ethics when they think the pay-offs will be high. This could make it *pro tanto* wrong to use the treatment.²⁰ To test the proposal, then, we should try to think of a case in which someone would be non-incidentally benefiting from wrongdoing and in which there are no other plausible candidate wrong-making features.

Consider:

The Inheritance: In his will, Clarence promises \$50,000 to a local theatre company, which he adores. It means a lot to him to know that this money will go to supporting them after he is gone. Vicious Cindy learns about Clarence's will. She has long been infatuated with the director of the theatre company, Joe, and she – delusionally – convinces herself that if she kills Clarence, getting the money to the company sooner, Joe will be so grateful that he will want to be with her. Cindy goes ahead with her plan. Her crime is discovered and she is convicted of murder. Everyone, including Joe, is suitably saddened and disgusted by Cindy's actions. Needless to say, Cindy does not get what she delusionally thought she would get from the murder. As per Clarence's wishes, \$50,000 of his inheritance go to the theatre company. And, because they can invest and use it sooner, they benefit much more than they would have had they not received this money until what would have been Clarence's natural death.

²⁰ Thanks to Wayne Sumner for a helpful discussion of this point.

Cindy has clearly acted wrongly in killing Clarence. And the beneficial consequences for the theatre company were not at all incidental: Cindy's wrongdoing was specifically intended to benefit the theatre company in this way, as a means to her own ends. Yet, there is nothing wrong (not even *pro tanto*) in the theatre company's following Clarence's wishes and accepting the money promised to them in his will. This shows that the fact that one would be knowingly benefiting from someone else's wrongdoing is not itself enough to make one's act wrong. Actions that involve benefiting from wrongdoing are, of course, sometimes wrong. But when they are, it is something else that makes them wrong.

3.2 Expressive Reasons

A different answer to the question of what moral point there is in refraining from eating meat in 'mere waste' cases like Leftovers is that our food choices have symbolic or expressive significance, independent of any causal influence they might have. By choosing not to eat Aunt Neta's leftover chicken, despite how tasty it looks, one expresses, even if only to oneself, one's commitment to animal welfare and disapproval of factory farming. The same goes in normal cases of deciding whether or not to purchase and eat meat. Even if your purchase of a chicken will not have any effect on what happens on the factory farm, by refraining you express your disapproval of the chicken industry and your concern for animal welfare. This – the claim goes – is something you ought to do. On the other hand, if you purchase factory-farmed meat, in the knowledge of the horrendous treatment of animals on these farms, this expresses support for that horrendous treatment, and this is something you ought, or have reason, not to do. This expressive proposal is the view that I will focus on most closely.

Now, one way the expressive power of an action can give you reason to do it is if, through this expressive power, you can influence others to act similarly. But this falls under the

“indirect difference-making” approach we already discussed.²¹ For the proposal here to be distinct, the idea has to be that what your act expresses matters morally *in itself*, even if it has no potential influence on the conduct of others. In other words, the reason has to be purely expressive.²²

One question is whether we really think there are moral reasons of a purely expressive variety. I won’t take a stand on this here, and will instead simply assume for the sake of discussion that there are. A second question is: even if we do have purely expressive moral reasons, are these going to be able to support any sort of moral requirement to avoid purchasing factory-farmed meat, to purchase fair trade coffee, etc.? Christopher Kutz, who appeals to purely expressive reasons, doesn’t think so, and concludes from this that they are only a partial solution to the problem. He writes:

It is doubtful whether any plausible construction of morality could require outcome-independent self-sacrifice [...] Like other supererogatory actions, symbolic actions fit awkwardly with a pure ethics of obligation. Nonetheless they have a firm place within a broader ethics of value.²³

However, if one really accepts the idea that what one’s act expresses matters morally in itself, it is not clear why one should rule out that there can be obligations coming purely from expressive considerations. In particular, at least insofar as the reasons you have to want to consume factory-

²¹ Lawford-Smith’s main proposal falls in this category. According to Lawford-Smith, one ought to buy ethically sourced products because this can be a way of signaling one’s willingness to form with others a collective agent that will work toward stopping the relevant harms or injustices. She claims that signaling one’s willingness to cooperate, through buying certain products rather than others, “prevents others from believing that their own cooperative actions would be futile, and sends the message that collective action is possible.” (2015, p. 322). This, she argues, is an important step toward actually forming such a collective. This proposal is a version of the indirect difference-making proposal. As Lawford-Smith says, it is “a consequentialist justification to actions that are normally taken to have mere expressive value”. (2015, p. 325). I think Lawford-Smith is pointing out something important in arguing that our purchasing-decisions can, in certain contexts, play this sort of role. But, for the same reasons as those discussed above in considering the indirect difference-making strategy, this proposal cannot satisfactorily solve the problem of collective impact.

²² The idea that our acts can have moral value purely in virtue of what they express has been advanced by Adams 1999 and Hill 1979. Kutz 2000 advocates for the expressive proposal as a partial solution to the problem of collective impact. (see Chapter 6). The expressive proposal also has been prominent as a reply in the case of voting and various other political activities (e.g. attending a large political protest). See, for instance, Benn 1979 and Brennan and Lomasky 1989.

²³ Kutz 2000, p. 191.

farmed meat, or drink unethically produced coffee, are relatively weak matters of personal preference (e.g. the taste, or wanting to spend less money so that you have more for entertainment), why not think that the expressive considerations outweigh those reasons, and thus obligate you? To use an example from Shafer-Landau, we might think that one acts wrongly if one voices “support for a racist dictator overseas”, and we might think this “even if one’s expression has no causal impact whatever on the treatment of the oppressed.”²⁴ Similarly, we might think one acts wrongly if one voices support for factory farming, even if one’s expression won’t have a causal impact.

But there is a much more decisive, internal problem for the expressive approach. The problem is: Even if it’s true that we have purely expressive moral reasons, these will only be reasons to do the things we are talking about – to refrain from eating factory-farmed meat, to choose fair trade, and so on – if those choices will have the expressive significance in question. But they do not necessarily have that significance. And when we think about what needs to be true for them to have it, we find that expressive reasons cannot be said to solve the problem of collective impact.

Consider the following two cases.

Lisa: Lisa is out for dinner. She orders the vegetable lasagna, rather than a meat option, because she feels like it, and not because of anything to do with animal welfare.

In this case, at least insofar as her reasons for choosing the lasagna are clear to her and those around her, Lisa’s choice does not seem to express concern for animals or disapproval of factory farming.

Ben: Ben is out for dinner. He orders the Cajun chicken rather than a vegetarian option because he loves the taste of Cajun chicken and because he justifiably and correctly believes that his ordering it won’t make any difference for the worse to animals. He cares a great deal about the

²⁴ Shafer-Landau 1994, p. 96.

wellbeing of animals, and if he had any reason to believe that his ordering this chicken could negatively affect animals, he would have chosen a vegetarian option.

At first pass, at least insofar as his motivations are clear to him and those around him, Ben's choice does not express disregard for animals or approval of factory farming. He has no such attitudes: he cares about animal welfare, and he disapproves of factory farming. And his choice does not express otherwise, because it is clear to everyone that he is only making it because he justifiably believes it won't have any negative impact on animals.²⁵

As we will see, these initial verdicts about Ben and Lisa could turn out to be false depending on certain further factors. But what these initial verdicts bring out is that just because a product was made in an unethical way, or by an industry that has harmful or unjust practices, this is not *by itself* enough to entail that your choice to consume it expresses an attitude of approval of those practices or of disregard for those who were harmed in the process. And it is not enough to entail that your choice to *refrain* from consuming it expresses an attitude of disapproval of those practices or concern for those who were harmed.

How could these choices gain the expressive significances in question? Here's one way: if there is an independent solution to the problem of collective impact, and we either know or should know it. By "independent solution" I mean a solution that does not appeal to expressive reasons – so, some other morally relevant consideration that explains why individuals should make the choices in question. Suppose, for the sake of illustration and contrary to what I said earlier, that the 'benefiting from wrongdoing' proposal is correct. So, suppose it's true that we ought not to eat factory-farmed meat because it is wrong to benefit from the abhorrent treatment of animals. If that is so, and if I know this or even just should know it, then if I go ahead and purchase a factory-farmed chicken anyway, this does express disregard for animals. It expresses

²⁵ See Driver 2015, pp. 74-75 for a related point.

disregard because in making this choice I am ignoring or failing to see the reason that their suffering gives me. Conversely, if you refuse to eat chicken because you think it is wrong to benefit from the wrongful infliction of suffering on animals, then your refusal expresses disapproval of factory farming practices.²⁶

If the food choices in question gain expressive significance via an independent solution, then expressive reasons can provide additional grounds for making the right choices. But, obviously, these expressive reasons would not be *solving* the problem of collective impact. The expressive reasons only apply here in virtue of the fact that there is a prior solution.²⁷

What this means is that to see if expressive considerations could solve the problem, we need to consider how the choices in question could carry the relevant expressive significances if there is no expressive-independent solution. So, suppose that if there is a moral story as to why individuals should refrain from purchasing factory-farmed meat, or buy fair trade coffee, and so on, it will have to be one that appeals to purely expressive considerations. All other options have been correctly ruled out.

Recall Ben: it seemed initially that he does not express any bad attitude toward animals when he chooses the chicken for dinner. He is only making this choice because he knows this won't make any difference for the worse to animals. We have already seen that this was too quick: if there is another independent moral story as to why he ought not to eat it, and one having to do with animal suffering on factory farms, and if he does or at least should know this, then his choice would still express disregard. But if there is no relevant non-expressive story as to why

²⁶ Note the asymmetry: with a known independent solution to the problem, the choice to go ahead and eat meat can express disregard, even if the person is not doing it for any 'anti-animal' reasons. In contrast, the choice to *refrain* from eating meat would not automatically carry the 'positive' expressive significance. For it to have that significance, you need – I think – to make the choice *for* those reasons, or at least under the appearance that those are your reasons.

²⁷ The same point holds if the choice gains expressive significance because we falsely believe there is an independent solution.

eating meat is wrong, and if everyone knows this, how can we say he is expressing disregard by making that choice?

There is one way: Ben's action could express disregard for animals if there is a societal convention that imbues this sort of choice with that expressive significance. Wearing green clothes on Earth Day can symbolize one's commitment to the environment. Wearing a shirt with the Confederate Flag on it expresses racism, or a lack of concern for those who have been victims of racism. These acts carry the expressive significance they do because of some process by which society imbued them with that meaning, and not through the fact that there is an independent moral story about why one ought to wear all green on April 22, or refrain from wearing a red shirt with a blue 'X' with white stars on it. Our individual food choices could acquire expressive significance in a similar way. Perhaps, then, the act of purchasing a chicken symbolizes *by way of societal convention*²⁸ a lack of concern for animal welfare, whereas purchasing lentils symbolizes the opposite.

If our food choices did carry expressive significance through societal convention, then if there are the purely expressive moral reasons in question, these will indeed be reasons to refrain from eating meat, to buy fair trade coffee, etc. But even if all of this is so, this does not solve our problem. The very fact that we can imagine a society in which the problem of collective impact arises, and in just the same way as it does in our society, but in which there are no such expressive conventions in place, shows this. The issue here is the same as with the indirect difference-making proposal. The expressive reasons would only apply in virtue of an empirical

²⁸ I am using the term "convention" only in a weak sense: just to mean that these choices were imbued, to some extent arbitrarily, with these meanings by some sort of social process. I do not mean to commit to any particular view about how something (an action, a word) can be imbued with meaning in a society. In this weak sense of "convention", all word-meaning is conventional, regardless of one's theory about how words acquire meaning. (See the introduction to section 7 of Rescorla 2015). Thank you to Olivia Sultanescu for drawing my attention to the need to clarify this.

contingency (in this case: a contingent expressive convention) that is not necessitated or made true by any of the conditions that give rise to the problem. Such reasons do not address the theoretical puzzle at issue. And practically-speaking: they do not give a stable, secure answer to the question of why individuals should refrain from making choices that will collectively result in tremendous harm and injustice. They just tell us why *insofar as* there are certain contingent conventions, individuals should refrain from such choices. That is a tenuous position to be in.²⁹

There is a final way one's choice could carry the relevant expressive significance: one could personally imbue it with that significance. I could declare, "I choose the vegetable lasagna to symbolize my disapproval of factory farming." But while I might have moral reason to express this attitude, and while this would be one way of doing it, in the absence of either an independent solution *or* a symbolic convention, there would be nothing morally requiring or recommending that I express it in that way. In other words, even if there is a moral reason to express my disapproval of factory farming, in the absence of an expressive convention or an independent solution, this would not give me reason *specifically* to refrain from eating meat. In such a world, I could just as well express my views by wearing a t-shirt with "Factory Farms are Evil!" printed on it. So, pointing out that I could imbue my choice with this significance cannot explain why individuals have moral reason specifically to make the food choices in question, and thus is a non-starter when it comes to the problem of collective impact.

²⁹ It's worth saying that I am doubtful that there are expressive conventions like these in our society at present. I think we can see this with the example of Ben. If there were a societal convention that makes *ordering chicken* an expression of disregard for animal suffering, then his "it won't make a difference to animal suffering" justification would do nothing in our ears to remove the expressive significance, just as someone who wears confederate-flag pajamas in the privacy of their own home cannot remove the expressive complaint in virtue of the fact that they are only doing so because they know that it won't make a difference for the worse to anyone. I don't mean to suggest that in our society, acts of purchasing meat do not express disregard for animal suffering, or that they do not do so for people with Ben's motivations. As explained, such acts could still express those attitudes if there is an expressive-independent solution to the problem, or even if some people just *believe* that there is. And the latter is surely the case. The point is just that these expressive significances don't seem to be coming from societal convention.

In summary: expressive considerations do not work to solve the problem. They might provide a moral reason or even a requirement to (specifically) make the choices in question, but whether they have the potential to do so depends on whether there is an independent solution to the problem of collective impact or a relevant sort of expressive convention. Either way, they cannot do the work of solving the problem.

4. Instrumental Contribution and Difference-Making

At the end of section 2 we took on board the conclusion that for many food choices, what one does individually won't make a difference with respect to the harms and injustices perpetrated by the food industry. From there, we turned to discuss some candidate 'non-instrumental' solutions. Our search there was not successful. In this final section, I will suggest that the reason for that lack of success is not that we did not look hard enough in that area, but rather that it is the wrong place to look.

Return to Leftovers. Does it make sense for a vegetarian to refuse on moral grounds, and specifically, grounds having to do with animal suffering, to eat meat in a case like this? I think it could. There are at least two explanations that could work. One is the temptation story mentioned earlier: if it is true of you that eating this chicken would weaken your resolve in 'normal' cases of deciding what to eat, this could give you moral reason, or a requirement, not to eat it. As we said, this only goes through if there is a prior solution to the problem of collective impact: an answer as to why you normally ought not to consume meat, even though it seems it won't make a difference. But if there is such a solution, this story can make sense.

The second story is an expressive one. I think it's easy – almost automatic – for a vegetarian to personally imbue her refusal in a case like Leftovers with expressive significance. She need not make any declaration like, "I am not going to eat this as a symbol of my

disapproval of factory farms”. Rather, she can simply not eat it *because* she is vegetarian. By not eating it because she is vegetarian, she makes her choice symbolic of her commitment to vegetarianism, and thus of the concerns that lead her to become vegetarian. So, if she became vegetarian because of concerns about animal suffering, her refusal to eat the meat in Leftovers expresses those concerns. If there are moral reasons to express these things, this would be a good way of responding to those reasons. But notice: this story also presupposes a prior solution to the problem of collective impact. It operates by way of the assumption that concerns about animal suffering support being vegetarian in the first place.³⁰

It thus seems that the best explanations for why it can make sense for a vegetarian to refuse, on moral grounds, to eat meat in cases like Leftovers involve presupposing that there is a solution to the problem of collective impact. When there is moral reason not to eat meat in ‘mere waste’ cases, this is often derivative on an answer as to why in regular cases one ought not to consume it. This suggests that we’re not going to solve the problem of collective impact by treating regular cases of purchasing factory-farmed meat, fair trade coffee, etc. – the cases in which the problem arises – as morally the same as ‘mere waste’ cases.

The temptation to treat them similarly comes from what the “it won’t make a difference” claim does. What this claim does, insofar as we accept it, is make ordinary cases of purchasing meat, conventional coffee, and so on, seem *instrumentally on par* with eating Aunt Neta’s chicken in Leftovers: insignificant or inert. If it won’t make a difference whether or not I buy meat – if things will be the same with respect to the harms in question either way – this seems to mean that my decision is instrumentally insignificant. It gives the impression that it won’t do any good for me to refrain, and that it won’t do any harm for me to go ahead, just like in ‘mere

³⁰ Note that neither of these two stories implies that a vegetarian who eats meat in ‘mere waste’ cases does anything wrong.

waste' cases. I don't think we can solve the problem of collective impact by accepting this picture and turning to non-instrumental considerations. Instead, we need to resist it. We need to show that your decision typically does matter instrumentally.³¹

This isn't to say that I think we can solve the problem, after all, by showing that one's individual purchases typically can make a difference. For the reasons discussed in section 2, I am doubtful that will work. In my view, the answer lies, rather, in seeing that one can make a significant instrumental contribution *even if* one cannot make a difference. If my purchasing a cup of fair trade coffee won't make a difference to the lives of coffee farmers, this tells me that no farmers' lives are going to go differently *depending* on whether I go with the fair trade option. But even if that is true, does it follow that in opting for the fair trade option I am not doing anything instrumentally useful toward improving the lives of coffee farmers? I think the answer is 'no'.

Notice that the solutions that we considered in section 3 did not really utilize an important feature of the problem of collective impact: the fact that by making certain food choices *going forward* we together could make a big difference for the better. I think a reply needs to capitalize on this fact. Now, consider: what does this fact tell us about what you individually can do?

In the case of meat consumption, if very many people refrain, going forward, from buying factory-farmed meat, this will force the meat industry to change or scale-back. What this tells us is that if *I* refrain from buying factory-farmed meat I will be doing precisely what many

³¹ One might think that another option is to identify a moral difference between normal cases and 'mere waste' cases, while accepting that they are instrumentally on par. Driver (2015) takes this approach, arguing that in normal cases one is complicit in wrongdoing if one purchases meat, *even if* one's purchase is instrumentally insignificant, while in mere waste cases one is not complicit. In Nefsky 2015, I discuss proposals of this sort, and argue that they cannot work as solutions to the problem. In particular, when it comes to complicity (or, participation in wrongdoing), I argue that for an appeal to complicity to address the problem, it needs to be that one is complicit in a sense that involves real instrumental contribution.

individuals must do if this possibility is to be actualized. That is, I will be doing something that *someone must do* – and someone else, and someone else, and so on – if this shift in consumer demand is to occur and thus, if this viable route to preventing grave harms is to be taken. Doing such a thing *is* making instrumental progress. You are not by yourself changing the outcome, but you are taking a significant step toward change.

I do not pretend, though, that it is easy to see how this amounts to making instrumental progress. Much more needs to be said to make this precise, and to debunk the impression that difference-making is required.³² In my view, the core answer as to why I have moral reason to make the food choices in question is the naïve, simple one: doing so makes real progress toward preventing grave harms or injustices. But what is hard is seeing *how* this could be true when one's doing so won't make a difference.

³² I develop this proposal in Nefsky 2016.

References

- Adams, Robert. 1999. *Finite and Infinite Goods: A Framework for Ethics*. New York: Oxford University Press.
- Anomaly, Jonathan. 2015. "What's Wrong With Factory Farming?" *Public Health Ethics* 8(3): 246-254.
- Benn, Stanley. 1979. "The Problematic Rationality of Political Participation," in *Philosophy, Politics, and Society: Fifth Series*, ed. Peter Laslett and James Fishkin. Oxford: Basil Blackwell.
- Brennan, Geoffrey and Loren Lomasky. 1989. "Large Numbers, Small Costs: The Uneasy Foundation of Democratic Rule," in *Politics and Process: New Essays in Democratic Thought*, ed. Geoffrey Brennan and Loren Lomasky. Cambridge: Cambridge University Press, 42-57.
- Budolfson, Mark. Forthcoming. "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequence Response," *Philosophical Studies*.
- Barry, Christian and David Wiens. Forthcoming. "Benefiting from Wrongdoing and Sustaining Wrongful Harm", *Journal of Moral Philosophy*.
- Driver, Julia. 2015. "Individual Consumption and Moral Complicity". In *The Moral Complexities of Eating Meat*. Ed. Ben Bramble and Bob Fischer. Oxford University Press, 67-79.
- Glover, Jonathan. 1975. "It Makes No Difference Whether or Not I Do It." *Proceedings of the Aristotelian Society* 49: 171-190.
- Hill, Thomas E. 1979. "Symbolic Protest and Calculated Silence." *Philosophy & Public Affairs* 9(1): 83-102.
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy & Public Affairs* 39(2): 105-141.

- Kutz, Christopher. 2000. *Complicity*. Cambridge: Cambridge University Press.
- Lawford-Smith, Holly. 2015. "Unethical Consumption and Obligations to Signal" *Ethics & International Affairs* 29 (3): 315-330.
- Nefsky, Julia. 2016. "How You Can Help, Without Making a Difference". *Philosophical Studies*. doi:10.1007/s11098-016-0808-y
- Nefsky, Julia. 2015. "Fairness, Participation, and the Real Problem of Collective Harm". *Oxford Studies in Normative Ethics* 5: 245-271.
- Nefsky, Julia. 2011. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy & Public Affairs* 39: 364-395.
- Norcross, Alastair. 2004. "Puppies, Pigs, and People: Eating Meat and Marginal Cases". *Philosophical Perspectives* 18: 229-245.
- Rescorla, Michael. 2015. "Convention", *The Stanford Encyclopedia of Philosophy*, Ed. Edward N. Zalta, <http://plato.stanford.edu/archives/sum2015/entries/convention/>.
- Shafer-Landau, Russ. 1994. Vegetarianism, Causation and Ethical Theory, *Public Affairs Quarterly*, 8(1): 85-100.
- Singer, Peter. 1980. "Utilitarianism and Vegetarianism". *Philosophy & Public Affairs* 9(4): 325-337.