King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Littlejohn, C. M. (2015). Potency and Permissibility. In B. Bramble, & B. Fischer (Eds.), *The Moral Complexities of Eating Meat* Oxford University Press, USA.

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Potency and Permissibility

Clayton Littlejohn

King's College London

Introduction

Smoking is pleasurable. It isn't just pleasurable; it is *immensely* pleasurable. A smoker can derive more pleasure from a single cigarette than, say, most of us could derive from eating a steak. Not only that, smoking is *distinctly* pleasurable. There is no good substitute for smoking a cigarette. When you want the pleasures associated with smoking a cigarette, only a cigarette will do. While pleasure is unarguably a contributor to your well-being (and arguably the only thing that contributes directly to your well-being), nobody would argue that you ought to smoke in order to maximize well-being. Isn't this paradoxical?

No, it isn't. It isn't even mildly puzzling. Even if hedonism is true and your concern is only with your own well-being, you'd have to take account of the pleasures you'd derive from smoking and then weigh them against the pleasures you could attain if you quit or decided against starting. Even if you don't get lung cancer from smoking, it can shorten your life or interfere with your ability to enjoy an active lifestyle. There's a very good chance it will harm you by causing you to undergo something rather unpleasant, but there's also a very good chance that it will harm you by robbing you of pleasures you could have experienced otherwise.

This isn't a paper about smoking. Philosophers don't write about smoking unless they're taking money from a tobacco company. Sadly, I'm not. This is a paper about the unreflective carnivore. Most of the people near and dear to me (myself included) are unreflective carnivores. We act as if we don't see any good reason not to order meat in restaurants or buy it in grocery stores. If you press us, we'll say that we eat meat because it brings us immense and distinctive pleasures. We seem to be under the impression that we maximize our own well-being by eating the significant quantities of meat that we do.

This kind of thinking is as confused as it is common. People who see that an argument from hedonism in favor of smoking is muddled should know better than to think that acting like the unreflective carnivore is a good strategy for maximizing well-being. The harms to health caused by consuming the amount of meat consumed by the typical meat consumers in the United States or Great Britain aren't negligible. Moreover, if we're honest with ourselves and we think about the amount of pleasure we derive from eating steak, we cannot really say that eating steak is *immensely* pleasurable. The pleasures associated with eating steak aren't like those associated with smoking much less like those associated with taking cocaine or ecstasy. My guess is that once you cultivate a taste for vegetarian cooking, you wouldn't rate your vegetarian meals lower than, say, your steak or ribs. The only reason people think otherwise is that they haven't really bothered to test this out for themselves.

This isn't a paper about whether cultivating a taste for vegetarian cooking is good strategy for maximizing well-being. It probably is, but I probably cannot convince you that it is unless you test my hypothesis by developing a taste for vegetarian cooking and do some research on the health benefits of a vegetarian lifestyle. This is a paper about the unreflective

carnivore, someone who makes decisions about what to purchase on the basis of a yen, a craving, or mistaken belief about what maximizes their own well-being. The question is whether there's any justification for acting like the unreflective carnivore.

If the unreflective carnivores near and dear to us are asked to reflect, they'll admit that they wouldn't set dogs on a rabbit and wouldn't set fire to a chicken. They wouldn't run a lamb down with their car. They wouldn't shoot a dolphin from the deck of a boat. When they watch nature shows, they are horrified if they see a baby elephant killed by lions and tear up when they watch the desperate mother elephant search for her young. In spite of this, they'll order rabbit, chicken, or lamb in a restaurant. They think that it matters, in some sense, what happens to animals, but they think that there's a sense in which their actions make of any moral significance. I want to see if these 'no difference' defenses hold up to scrutiny.

It's important to fix some parameters for this discussion. The people near and dear to me live in the United States and Great Britain. They can afford to eat in restaurants and purchase meat in the grocery store. They can also afford to eat fresh fruits and vegetables, which happen to be in abundant supply where they live. If anything, they'd save money by becoming vegetarians. Although such a diet isn't appealing to them, they know enough about cooking that they could easily learn to cook delicious and healthy vegetarian meals if pushed to do so. They don't suffer from any bizarre medical conditions that would require a daily dose of bacon. They aren't utility monsters. They don't derive freakish amounts of pleasure from eating steak or chicken potpie. When I've served them vegetarian meals, they haven't rated them lower than they'd rate a chicken or pork dish that I'd prepare or that they'd prepare. They know that the animals they eat are sentient, capable of feeling pleasure and pain. They believe that some of the animals they eat lived reasonably good lives and were killed humanely. They believe that

some of the animals they eat lack complex mental lives and so cannot have self-regarding desires about their own futures. They also realize that the animals they eat are typically killed in vast numbers so that the decision to kill a chicken or lamb, say, isn't directly in response to the order they place in a restaurant or a purchase they make in a market. This should give you some sense of the kind of person I have in mind when I evaluate the no difference defenses.

Let me fix one more parameter. The people near and dear to me aren't neo-Cartesians and so they agree that some of the animals they eat are capable of suffering as well as capable of enjoying pleasurable experiences. They agree that this means that some of the non-human animals we eat are moral *patients*, creatures who have interests that have non-instrumental moral significance. They might not describe their view in these terms, but this is a view that they're committed to and something they'll accept without too much of a fuss.

To see why they should accept this, let's review the marginal cases argument.¹ The starting point of the argument is this:

MC1: All living, sentient humans are moral patients.

That is to say, all living, sentient humans have interests that carry some sort of moral weight. It can be *prima facie* wrong to do something that harms those interests or fails to promote those interests. The next step is this:

MC2: If all living, sentient humans are moral patients, each of these individuals must possess some set of features that's sufficient for conferring the status of moral patient upon them.

The idea here is that for any being that has the status of being a moral patient, there is some feature or non-empty set of features that's sufficient for conferring that status upon them. Those who accept (MC2) insist that there's some reason why moral patients have that status and that the possession of this status will be connected to the features of the individuals that possess it. If you accept this, then it's hard to see how you could avoid saying the following:

MC3: Any set of features that's sufficient for conferring the status of moral patient upon the living, sentient human animals that have them would be sufficient for conferring the status of moral patient on any non-human animal that has this set of features.

Moreover, it seems that the following is quite plausible:

MC4: There are non-human animals that we eat that have these features.

If this is correct, we have our conclusion:

MCC: There are non-human animals that we eat that are moral patients.

Once this conclusion has been established, we have to acknowledge that there are non-human animals that have a moral status that ground obligations. If you're a consequentialist, you'd have to acknowledge that the considerations about aggregate well-being that determine the rightness of your actions will include considerations about the well-being of some non-human animals. If you're a non-consequentialist, you'd have to acknowledge that there are principle-based protections that apply to non-human animals that are similar to those that we recognize for humans at the margins of life. It will matter later that the argument doesn't simply establish that animals matter morally, but that they matter morally in particular ways that humans do.

If readers don't want to accept the argument's conclusion, so be it. Objections to this argument have been dealt with elsewhere and my unreflective carnivores don't try to undermine this argument to defend their behavior.² Still, a brief word about (MC4) is in order. In offering this argument, I didn't say that there aren't significant moral differences between puppies and children. There might well be significant moral differences between all living, sentient humans and all non-human animals. That doesn't matter. All that matters is whether these differences are sufficient to undermine the argument that there are fish, chickens, cows, ducks, rabbits, pigs, lambs, cows, buffalo, etc. have *some* kind of moral status. Think about human infants and three ways that a human infant might be:

- a. Sentient, capable of having self-regarding desires about its future, capable of developing into an adult with normal intellectual and emotional capacities.
- b. Sentient, capable of having some self-regarding desires about its future, but incapable of developing into an adult with normal intellectual and emotional capacities.

- c. Sentient, incapable of having some self-regarding desires about its future, and incapable of developing into an adult with normal intellectual and emotional capacities.

If you thought that only infants with the characteristics described in (a) had moral status, you could try to resist the argument for (MCC) on the grounds that you thought that the non-human animals we eat don't have the features necessary for moral status. The trouble with your position is that you could only deny (MC4) by excluding many infants from the moral sphere. That's monstrous. On the view you'd be defending, there would be healthy infants that weren't moral patients. Since they wouldn't have any interests that carried any moral weight, your view would be that there wouldn't be any principled objection to eating them or using them for parts. (At least, there wouldn't be any such objection that appealed to the status possessed by such infants since your view would be that they didn't have any.) If, on the other hand, you stand rightly opposed to eating all human infants by virtue of the fact that all of them have moral status, you couldn't resist (MCC) unless you thought that the non-human animals we ate weren't sentient. You don't think that. Insofar as you think that sentience is sufficient to confer moral status upon all sentient infants, you should think that it is sufficient to confer a kind of moral status upon the sentient animals we're eating.

<1>

If the marginal cases argument is sound, you shouldn't be terribly impressed with our first no-difference argument:

I'd agree that you shouldn't eat things like monkeys, dolphins, or dogs because they have the higher-capacities needed to form self-regarding desires about their futures. When we're dealing with creatures like that, we're dealing with creatures that can be harmed by being killed because their futures can be good for them. When we're dealing with simpler sentient creatures incapable of thinking about themselves in the future, however, their deaths cannot be bad for them because they cannot have any self-regarding desires about what happens to them in the future. So long as they're humanely housed and killed, no wrong has been committed. Thus, it makes no moral difference whether we eat sentient creatures or not, provided that they lack high-order cognitive capacities involved in forming self-regarding desires and they were killed humanely.

The argument isn't the slightest bit convincing when it's applied to infants that lack the potential for developing the capacities necessary for having self-regarding desires about their own futures, so it shouldn't be any more convincing when it's applied to creatures like fish or chickens.³ On the most plausible account of why our deaths are bad for us, deaths are bad for us by virtue of depriving us of something valuable; death is bad by virtue of the fact that it deprives us of a future period of existence that would have contained a positive level of well-being. This account explains why the death of disabled infants can be bad for the infant and it applies to sentient non-human animals, too.⁴

Before pressing on, there is a point that I should clarify. I don't think that *this* is how we should response to the first no difference argument:

Death's badness is due to the fact that it deprives something of a future in which it is a sentient creature that lives a life that is overall good for it. Thus, it is wrong to kill sentient animals if they could have continued living lives that were overall good for them.

This response moves too quickly from a claim about the kinds of events that can harm an individual to a claim about what sorts of events it would be wrong to cause. This response assumes, implicitly, that causing a deprivation involves *harming* and *wronging* and I think we should leave it open whether harms invariably wrong.⁵ The deprivation account of death's *badness* seems to imply that an earlier death is typically going to be a worse death, at least if we're assuming that the relevant individual that stands to live a life that's overall and uniformly good. If the *badness* of a death (which is measured in terms of the magnitude of a deprivation) is a sufficient ground for the *wrongness* of causing death, the deprivation account of death's badness implies that it is wrong to terminate a pregnancy when bringing the fetus to term would have resulted in a life that was overall good for the individual. Indeed, it seems to imply that there's a wrong in the case of abortion and infanticide that's roughly equivalent.

The way to avoid this implication is to draw a distinction between the deprivations that cause harm or leave someone worse off and the deprivations that do so in such a way that's wrongful. If we assume that all deprivations are harms, we shouldn't assume that it is invariably wrong to cause them. (If we assume that it's always wrong to harm, we shouldn't assume that all deprivations harm. I'd prefer to speak as if deprivation and harm are the pair that's more intimately connected.) Here's one reason to draw this distinction. If we think that an infant once was a fetus (i.e., that the patient that is the infant was once identical to a fetus), we'd have

to say that terminating a pregnancy causes a deprivation to the relevant patient that's equivalent in terms of the magnitude of harm to the deprivation caused by an act of infanticide. And if we assume that the badness of the deprivation is itself the sufficient ground of the wrong of infanticide, it seems we get the result that aborting a pregnancy is no less wrong than infanticide because it's no less worse for the relevant individual.

It might initially seem odd that there's a gap between causing a deprivation/harm on the one hand and wrongdoing on the other, but examples can help to make this intuitive. Suppose we were to discover that planaria had some minimal degree of sentience. You'll recall that if you divide a planarian into two parts, each part will regenerate the parts needed so the result will be two living flatworms, not one dead flatworm in two pieces. Phyllis the planarian is about to be cut into two pieces but you intervene so she suffers only a slight nick, not a complete separation. She heals so we don't have a case with one dead planarian or two living planaria. I don't think that the parts of Priscilla that could have grown into full planarian had an interest in being split off, but there would be two individual planarian that would have been better off as a result.⁶

Of course, someone might ask at this point why we should think that the animals that cannot anticipate their own futures are patients in the sense that they have interests that carry moral weight. The answer is easy. Although there's a potential gap between harming and wronging, we don't think that the possibility of such a gap allows us to terminate the life of an infant in groups (b) or (c). The marginal cases argument gives us some reason to think, then, that the mental capacities common to these infants and the relevant animals are sufficient for conferring upon both the kind of moral status that makes it seriously wrong for us to harm them by killing them.

To understand the moral distinctions between non-creation, abortion, and infanticide, we shouldn't just look to considerations having to do with the aggregate well-being of non-creation, abortion, infanticide, and the relevant feasible alternatives. We can distinguish abortion from infanticide (and contraception) by pointing out that it's only in one of these cases that we're dealing with a moral patient. It isn't in the interests of the fetus to develop into an infant, you could say, just as it isn't in the interests of an egg or sperm to fuse with the other. The account that I'd recommend combines a deprivation account of harming with a patient-affecting restriction on wrongdoing. It should be noted that the account is incompatible with any consequentialist approach to right action that accepts totalism:

Totalism: The morally relevant outcome of an action is the possible world that would be actual if the action were performed.⁷

If you accept totalism, you wouldn't think that there would be much difference between deciding against bringing a patient into being and snuffing the patient out. This, however, is an important distinction if you accept the patient-affecting restriction since the restriction says in effect that only the snuffing is a violation of the principle of non-maleficence. There might be consequentialists that reject totalism, but it's hard to see how they could defend such a move. If you think that the good is prior to the right, it's hard to see how you could sensibly distinguish the relevant goods from the irrelevant goods by appeal to considerations like whether the relevant goods would be enjoyed by actual agents as opposed to merely possible agents you could bring into being. If we think of morality as primarily concerned with how patients are treated and only derivatively concerned with things like aggregate levels of well-being, the restriction might seem intuitive and consequentialism should appear to be problematic.

<2>

The first no difference argument fails because it tries to exclude certain sentient creatures from the moral sphere. The marginal cases argument should remind us of two things. The first is that when you raise the bar, you run the risk of excluding some living humans from the moral sphere. The second is that if we draw fine-grained distinctions between kinds of moral status, we'll find that some sentient non-human animals will have the same kind of moral status as some living humans. We shouldn't just think that animals are potential victims. We should also recognize that this potential for victimhood grounds obligations concerning animals that are similar to the obligations we have concerning fellow humans.

The second no-difference argument is supposed to avoid these worries because it doesn't aim to show that there are sentient animals that lack moral status. Instead, it's supposed to show that the unreflective carnivores often don't have the obligation to refrain from buying or ordering meat because of the causal structure of the supply chain:

In sitting down at a restaurant that has lobster on the menu, it doesn't matter what I order because if I don't order the last lobster else will.⁸ Thus, while there might be obligations not to harm animals, the existence of such obligations has nothing to do with me on this occasion.

The argument has limited import, of course, because it only applies to those cases where a living animal would be killed if someone were to request it. We'll discuss cases where the animals have already been slaughtered and prepared for consumption in the next section, so let's just focus on cases with this structure for now.

If the argument shows that the unreflective carnivore has no responsibility to refrain from eating the lobster, it also seems to show that the unreflective carnivore cannot be held responsible for deciding to order the lobster. Although it's possible to fail to meet your responsibilities without being responsible for the failure, there are nevertheless principled connections between responsibility in its forward- and backward-looking sense. If you *know* that you're under no responsibility to X, you cannot be held responsible for X-ing. Thus, if the argument shows that the unreflective carnivore knows that she's under no obligation to refrain, it also shows that the agent isn't culpable for her decision.

These results strike me as problematic. Even if the unreflective carnivore knows that another customer would order the lobster if she didn't, she would still be under an obligation not to order it. (Assuming, of course, that she'd be under this obligation if she were the restaurant's only patron.) If you like slogans, here's one: preemption doesn't bear on permissibility. Moreover, even if the unreflective carnivore knows that the lobster is going to get it one way or the other, she's still culpable if *she* is the patron who orders it. The question is how can we make sense of this in light of the fact that the lobster won't be made worse off than she would be otherwise by the agent's decision to order her for dinner.

Let's start with something simple. Suppose you think that there's a principle of non-maleficence, one that says that it would be *pro tanto* wrong for you to harm a moral patient.⁹ If you order the lobster and preempt another patron waiting in the wings, you cause the lobster's death (with help from the restaurant staff, of course). So, *you* harm the lobster. So, you commit the wrong, right?

One problem with this line of thought is that it seems susceptible to the following line of objection:

Look, you couldn't have harmed the lobster if the lobster isn't harmed as a result of your action. For the lobster to be harmed as a result of your action, it has to be that the lobster suffers a harm. This requires that the lobster is worse-off than she would have been otherwise. But the lobster isn't worse off than she would have been otherwise as a result of your action. If you didn't order the lobster, someone else would have. Since there was no harm, there wasn't any violation of the principle of non-maleficence.

It seems rather plausible that an agent doesn't harm a patient in our example if the patient doesn't suffer any harm by being killed. I certainly don't want to defend the view that you can harm the lobster even when the lobster isn't harmed. What should we say in response?

The objection assumes that harm should be understood *comparatively* in accordance with this counterfactual account of harm: an event is harmful for an individual iff things go worse for this individual than they would have gone if the event had not occurred (Bradley 2012: 396).¹⁰ Applied to the case at hand, you don't harm the lobster by ordering it and thereby ordering its death because if you hadn't done this, somebody else would have. Thus, someone who accepts this account can only say that you harmed the lobster if they were to say that it's possible for you to harm the lobster when the lobster doesn't suffer any harm at all. It would be best not to say that, I think.

Comparative accounts of harm tell us that a victim is harmed iff something makes the victim worse-off.¹¹ Victims cannot be worse-off *simpliciter*. Something harms a victim only if it makes the victim worse-off than the victim is made to be under some relevant alternative. It's quite natural to think that the relevant alternative to consider in determining whether a happening harms someone is the one in which the happening doesn't happen because when we are trying to decide what to try to avoid or prevent, we run (if we're rational) a kind of counterfactual comparison test. The lobster has little reason to do what it can to get you to act otherwise if she knows that another patron will order her for dinner if you don't. The important point here is that it would be strange for the lobster to think that she won't be harmed by any of the patrons in the restaurant because so many of them want her boiled alive.

Consider an alternative to the counterfactual comparison account. The revised comparative account of harm tells us to compare an occurrence to its non-occurrence to determine whether the occurrence harmed the victim rather than comparing what happened to what would have happened if the occurrence hadn't occurred. Some events are intrinsically bad for a victim, cause something that's intrinsically bad for a victim, or prevent the victim from enjoying something that's intrinsically good.¹² An event harms a victim iff it meets this disjunctive condition and the occurrence (along with its effects) are worse for the victim than the non-occurrence's effects would have been. This account is tailored to vindicate the intuition that even if there are multiple patrons who will order the lobster to be boiled alive, the particular patron that calls in the order is the one that harms the lobster. The order initiates a causal chain that results in the lobster being boiled alive and then dying. This is a worse state than the state that the non-occurrence of the order would cause. The non-occurrence wouldn't

cause the lobster to be in any state at all. This *isn't* because non-occurrences don't cause. It seems plausible that non-occurrences can cause (e.g., the failure to water a plant or feed a pet can cause death). The point is that this particular non-occurrence didn't cause the lobster to boil or to die and this allows us to attribute responsibility for the bad state the lobster is in to one order in particular.¹³

I don't think that consequentialists will like this account of harm. It's hard to see how a consequentialist could think that the concept of harm as its understood here could be a useful concept for the purposes of moral theory since it doesn't have much to do with the kinds of counterfactual assessments that allow us to determine which of the available alternatives will involve the greatest amount of aggregate well-being. I've suggested that morally conscientious agents should be concerned about the particular role they play in harming and that this requires us to think about causal relations between themselves and the moral patients they causally interact with. Consequentialists will likely say that it would be odd for the morally conscientious agent to be concerned with this. They would say that we ought to be concerned with aggregates, not our role in how these aggregate levels are produced.

One of Williams' (1988) examples might help to illustrate the consequentialist's concern. You are offered a choice. You can pick up a rifle and shoot one villager held by the general's men or do nothing. If you do nothing, the general's men will execute this villager and nineteen others. You might think (as Williams presumably didn't) that it would be perverse for you to think that your role in the death of a villager as the sort of thing that should enter your deliberation in this sort of case. I agree. It would be perverse to do so when it's so clear that

concerns about aggregate well-being override concerns you might have about your role in the killing of one or more villagers.

Having said that, I think we're on pretty good ground here. First, the view that I'd defend isn't a view on which considerations about the role you play in harming, say, override considerations about overall levels of utility. I'd say that there's a *pro tanto* reason for you to refrain from shooting a villager and that this reason can potentially be overridden if you have a way of shooting this villager so as to save nineteen lives. As this is just what the consequentialist wants us to say, this sort of case needn't cause us any serious concern. Second, we've already seen that our concern should be directed primarily towards patients and their well-being and so only indirectly concerned with aggregates. If the consequentialist case against the kind of harm principle I'm offering rests on the idea that our concern as morally conscientious agents should be with aggregate levels of well-being, the force the objection derives from an approach we've already seen to be problematic insofar as it lacks the resources to explain why we should draw a moral distinction between promoting the interests of actual patients and bringing patients into being so as to promote their interests. Third, it would seem that the right way to test the account on offer is to consider a modified version of Williams' case. If the choice we're offered is one in which we can decide to shoot or refrain from shooting and that same villager will be shot by the general's men, I'd say that there's a duty to refrain. If you share that intuition, it counts in favor of the view I'm recommending.

Finally, I should note that the account of harm offered here doesn't vindicate the idea that when you harm some victim, the victim is worse-off overall as a result. Some see this as a problem because they think that *to be harmed*, one has to be worse-off overall as a result and

think that if you harm someone, someone has to be harmed.¹⁴ While I agree that if you harm someone, someone has to be harmed, I don't think that to be harmed, you have to be worse-off overall as a result. Consider Ross' principle of non-maleficence, a principle that says that there's a *prima facie* duty not to harm. If we said that we had only a *prima facie* duty not to perform actions that leave someone worse-off overall, we'd lose sight of the fact that any of our actions that cause something intrinsically bad for a moral patient calls for a justification and is *prima facie* wrong. It's true that considerations about overall well-being *might* provide that justification, but that's consistent with the idea that every harm (as I've characterized it) might be *prima facie* wrong to bring about. Thus, in moral contexts, at least, it makes sense to characterize harm as I have, provided that we understand that it's one moral notion amongst many and shouldn't be considered to the exclusion of considerations about whether a victim might be better off overall by being harmed than not. I suspect that part of the intuition that harm should have to do with overall well-being has to do with the fact that in many non-moral contexts where we're thinking about which option would be best for us, we rightly focus on aggregate well-being rather than particular intrinsically good or bad events. I'd say that our focus in these contexts *should* be on overall notions and that this just goes to show that the notion of harm isn't nearly as useful in such contexts as the distinct notion of what's better for us overall.

<3>

A final challenge remains. Consider one final no-difference argument:

I can see that I might bear a special kind of responsibility for the causal contribution that I make, but the situation I face in a grocery store is unlike the preemption cases because my decisions in this setting don't make a causal difference to the well-being of any animal. The animals I buy were already dead when I arrived at the store and the decision to slaughter more animals for food isn't sensitive to my decision about whether to buy chicken, lamb, beef, etc. Since my actions don't have any causal impact on any living animals, I don't have to take account of considerations of animal welfare in deciding what to buy in the store.

Does this argument get the carnivore off the hook?

These arguments concern cases that might appear to be different in certain respects. The second no-difference argument focuses on preemption cases in which a particular patron's decision has the potential to cause the death of a moral patient. In this case, the consumer is trying to decide whether to purchase an animal that has already been killed. Moreover, the patron in the restaurant has very good reason to think that her decision can cause the death of the lobster. The consumer in the shop has very good reason to think that her decision won't harm any animal at all. She might say the following:

[I]f I did not buy and consume factory-raised meat, no animals would be spared lives of misery. Agribusiness is much too large to respond to the behavior of one consumer.

Therefore I cannot prevent the suffering of any animals. I may well regret the suffering inflicted on animals for the sake of human enjoyment. I may even agree that the human

enjoyment doesn't justify the suffering. However, since the animals will suffer no matter what I do, I may as well enjoy the taste of their flesh (Norcross 2004: 231).

Isn't she right?

This is the infamous causal impotence defense. There's a familiar consequentialist response to this argument, one that strikes me as being decisive.¹⁵ It's true that agribusiness is too large to respond to the behavior of just one consumer, but it has to respond to some consumer behavior. Let's say that every month the factory gets data about the number of chickens sold and will change its production in response to increases or decreases in demand. If, say, demand decreases by 1000, 1000 fewer chickens will be produced. Of course, the consumer might have good reason to think that her decision not to purchase won't result in a reduction of the number of chickens killed, but why should that matter? If you think that the permissibility of an action is determined by *expected* utility rather than actual utility, this is irrelevant. It's been conceded that the utility gained by killing and eating a chicken is less than the amount of disutility involved in killing the chicken. Once that's conceded, it also has to be conceded that the expected utility of buying chicken under the conditions described is also negative. There is one chance in a thousand that the consumer purchases the threshold chicken. If purchased, it triggers the death of 1000 chickens, which is a very bad result. If not, it avoids triggering the death of 1000 chickens, which is a much less bad result. And in terms of expected utility, the expected utility of taking the chance is the same as the expected utility of raising a chicken in your own little factor farm and killing it yourself.¹⁶

Of course, the consumer might think to herself that it really doesn't matter whether *she* buys the threshold chicken or not because somebody is going to buy it. Perhaps, but then this consumer is relying on the arms dealer defense considered above. If it doesn't work for war criminals, it shouldn't work unreflective carnivores. Imagine we gather all the consumers who purchased chicken in a particular grocery store and arrange them in line on the basis of their purchase. Every consumer, let's say, buys just one chicken. Suppose you are the 7th person in this line. You look back and see the consumer who purchased the 1000th chicken. She triggered a slaughter of 1000 more chickens, but she wouldn't have if you didn't do your part. This was a team effort. You could have chosen not to make your own contribution, which causally is no different from hers, but you didn't and so you share equal causal responsibility for the deaths just as you would if the whole lot of you grabbed a rope and pulled a very heavy cage filled with these chickens over the side of a cliff.

Suppose, however, that you look back down the line and the people after you and it just so happens that nobody bought the threshold chicken. As a matter of luck you aren't part of a unit that worked together to kill off 1000 chickens. Maybe the factory closed its shutters or the grocery store burned to the ground. If you think you are under an obligation not to impose risk, this just shows that your actions didn't harm, not that you didn't commit a wrong. Compare this case to the case where you play a round of Russian roulette and live to tell the tale. No harm came of what you did, but you shouldn't have done it. Moreover, let's not forget that the scope of your obligation concerning animals can be quite broad. Your financial support sustains the agribusiness outfits that produced your chicken. Even if your previous supplier closes shop,

your financial support sends signals to other producers who compete for your business. If we focus on just the factories that you buy from, we lose sight of this.

We also shouldn't lose sight of the fact that your behavior as a consumer helps to perpetuate the idea that it's acceptable for people to treat moral patients as resources to be used for making a profit with little regard to their welfare. Their behavior can be modified by public sentiment, but your buying behavior doesn't just encourage the further slaughter of animals, it helps to shield them from the moral judgments that could be efficacious in reforming their practices if only more people like you saw the light. Here's a conjecture: if a critical mass of people came to believe that it's morally repugnant for people to profit from an industry that causes widespread animal suffering, there would be fewer people who felt free to invest in it. If there were a social stigma attached to factory farming that's akin to the social stigma that's attached to, say, pornography or the exploitation of sweatshop labor, there would be fewer factory farms. Although you cannot choose whether there will be a critical mass of condemnation that would influence the behavior of investors, you can directly determine whether you'll be part of the critical mass that sustains a morally repugnant practice. As we've seen earlier, this is sufficient to show that you're under some obligation to refrain because it means that you're a part cause of a state that impedes a triggering event that results in a significant decrease in the amount of animal suffering.

Even if you are reasonably convinced that your own actions won't help to alleviate animal suffering, there's one further consideration to bear in mind. The ground of your obligation to refrain from eating animals has to do with animal suffering. The scope of your obligation can be quite broad. It needn't be limited to an obligation to take steps to reduce animal suffering. It might require you to perform actions that have no causal connection to the

welfare of any animals. If there are institutions or practices responsible for widespread suffering, you might be obligated to break all association with them even if doing so has no real causal impact on their operations. As social creatures, having a good social standing, one that's free from sanction and blame, might itself be a good, a good that's not appropriately enjoyed by those responsible for widespread death and suffering. If there were institutions responsible for the widespread death or the widespread abuse of children, for example, you wouldn't think that it would be acceptable for you to continue with any sort of involvement with these institutions by investing in them, working for them, or purchasing their products. The fact that we're comfortable sustaining connections with the institutions responsible for the widespread slaughter of moral patients suggests that we haven't yet internalized the obvious consequences of the marginal cases argument.¹⁷

Notes

¹ See Singer (1993) for an early discussion of the argument.

² See Kazez (2010) and Norcross (2004) for discussion of objections to the marginal cases argument.

³ Assuming, of course, that chickens and the relevant fish are sentient. For further discussion of deprivation and desire-based accounts of well-being, see Bradley (2009).

⁴ It's an interesting question whether non-culpable ignorance of the fact that sentience is sufficient for moral status can subvert obligation or exculpate. For arguments that it cannot, see Arpaly (2004), Harman (2011), and Littlejohn (2014). For arguments that it can, see Zimmerman (2008).

⁵ This kind of distinction is crucial for understanding Thomson's (1971) defense of abortion and nicely explains why the austere basis that Marquis (1989) tries to argue from don't support his conclusions about abortion.

⁶ For the purposes of the example, I'm assuming that if there's an individual that has some positive level of well-being (e.g., a sentient planarian that results from splitting Priscilla into two pieces) we can say that this individual would have been harmed if we had done something that resulted in a situation in which it isn't true that that individual exists and enjoys at least as much well-being. I appreciate that this is controversial, but I don't think anything I defend here turns on this. The example is a helpful reminder of the gap between deprivations and wrongs. See Kamm (2005) for further discussion of this point in connection with recent debates about the use of embryos in medical research.

⁷ See Carlson (1995: 10).

⁸ Gardner (2007: 72) dubs this kind of argument the 'arms dealer' argument.

⁹ See Ross (1930). Dancy (1998) provides a helpful discussion of the differences between Ross' understanding of harm and the duty of non-maleficence and the consequentialist understanding of the duty.

¹⁰ There are non-comparative accounts of harm to consider. See Harman (2009) for a helpful discussion of these approaches. While I think her approach is intuitively quite plausible, I don't think we have to accept a non-comparative account to deal with this particular objection.

¹¹ See Bradley (2009) for a defense.

¹² This is Bradley's (2009) notion of a *prima facie* harm. A complete account of harm should say something about *preventions* whereas the account that I've offered above doesn't. We should

compare the goods that an occurrence prevents to the goods that the non-occurrence would have prevented. That's unwieldy, however, so I've simplified the account above accordingly.

¹³ The modified comparative account of harm applies some ideas of Sartorio's (2005, MS) about causation and responsibility to deal with problems for comparative accounts of harm. To my knowledge, she hasn't tried to use her account of causation and difference-making to deal with problems that arise for the counterfactual comparison account, but it seems like a natural application. If it's helpful, credit clearly goes to her. If it isn't, credit will have to fall to me. For a critical discussion of Sartorio's approach, see Driver (2007).

¹⁴ Thanks to Ben Bramble for raising a version of this objection.

¹⁵ I believe this response first appears in Almeida and Bernstein (2000). For further discussion of consequentialist responses to the threshold argument, see Almassi (2011), Chartier (2006), Harris and Galvin (2012), Matheny (2002), Norcross (2004), and Kagan (2011).

¹⁶ For arguments that overall obligation depends upon expected value rather than actual value, see Jackson (1991) and Zimmerman (2008).

¹⁷ I want to thank Ben Bramble and Bob Fischer for their helpful comments on an earlier draft of this paper. Thanks also to Matt Burstein, John Harris, Robert Howell, Alastair Norcross, and Amy Revier for discussing these issues with me.

References

Almassi, B. 2011. The Consequences of Individual Consumption: A Defense of Threshold Arguments for Vegetarianism and Consumer Ethics. *Journal of Applied Philosophy* 28: 396-411.

- Almeida, M. and Bernstein, M. 2000. Opportunistic Carnivorism. *Journal of Applied Philosophy* 17: 205–11.
- Arpaly, N. 2004. *Unprincipled Virtue*. Oxford University Press.
- Bentham, J. 2010. *The Principles of Morals and Legislation*. Nabu Press.
- Bradley, B. 2009. *Well-Being and Death*. Oxford University Press.
- _____. 2012. Doing Away with Harm. *Philosophy and Phenomenological Research* 85: 390-412.
- Carlson, E. 1995. *Consequentialism Reconsidered*. Kluwer.
- Chartier, G. 2006. The Threshold Argument Against Consumer Meat Purchases. *Journal of Social Philosophy* 37: 233-49.
- Dancy, J. 1998. Wiggins and Ross. *Utilitas* 10: 281-85.
- Driver, J. 2007. Attributions of Causation and Moral Responsibility. In W. Sinnott-Armstrong (ed.), *Moral Psychology*, Vol. II. MIT University Press, pp. 423-39.
- Harman, E. 2009. Harming as Causing Harm. In M. Roberts and D. Wasserman (ed.), *Harming Future Persons*. Springer, pp. 137-54.
- _____. 2011. Does Moral Ignorance Exculpate? *Ratio* 24: 443-68.
- Harris, J. and R. Galvin. 2012. ‘Pass the Cocoamone, Please’: Causal Impotence, Opportunistic Vegetarianism, and Act-Utilitarianism. *Ethics, Policy, & Environment* 15: 368-83.
- Jackson, Frank. 1991. Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics* 101: 461-82.
- Gardner, J. 2007. *Offences and Defenses*. Oxford University Press.
- Kagan, S. 2011. Do I Make a Difference? *Philosophy and Public Affairs* 39: 105-41.
- Kamm, F. 2005. Ethical Issues in Using and not Using Embryonic Stem Cells. *Stem Cell Reviews* 1: 325-30.

- Kazez, J. 2010. *Animalkind: What We Owe to Animals*. Blackwell-Wiley.
- Littlejohn, C. 2014. The Unity of Reason. In C. Littlejohn and J. Turri (ed.), *Epistemic Norms*. Oxford University Press, pp. 135-54.
- Marquis, D. 1989. Why Abortion is Immoral. *Journal of Philosophy* 86: 183-202.
- Matheny, G. 2002. Expected Utility, Contributory Causation, and Vegetarianism. *Journal of Applied Philosophy* 19: 293-97.
- Norcross, A. 2004. Puppies, Pigs, and People: Eating Meat and Marginal Cases. *Philosophical Perspectives* 18: 229-45.
- _____. 2005. Harming in Context. *Philosophical Studies* 123: 149-73.
- Ross, D. 1930. *The Right and the Good*. Oxford University Press.
- Sartorio, C. 2005. Causes as Difference Makers. *Philosophical Studies* 123: 71-96.
- _____. MS. What Difference Does it Make? On Acting Freely and Making a Difference.
- Singer, P. 1993. *Practical Ethics*. Cambridge University Press.
- Thomson, J. 1971. A Defense of Abortion. *Philosophy and Public Affairs* 1: 67-95.
- Williams, B. 1988. Consequentialism and Integrity. In. S. Scheffler (ed.), *Consequentialism and its Critics*. Oxford University Press, pp. 20-50.
- Zimmerman, M. 2008. *Living with Uncertainty*. Cambridge University Press.