

Non-Ideal Food Choices

Tyler Doggett and Andy Egan

Our lives are full of failures. Some are significant. Some are trivial. Some are reprehensible. Some are laughable. Some are worth responding to. Some aren't. But, as we said, life is full of these things, so let's be gentle about them. Failing is a way of acting non-ideally. Instead of "failure" let's just talk about behaving in a "non-ideal" way. When the New England Patriots won the Super Bowl in 2015, this was, for the Patriots and their fans, ideal. When the New England Patriots won the Super Bowl in 2015, this was, for the fans of the Cleveland Browns, non-ideal: There can be only one Super Bowl winner and the Browns failed in pursuit of the ideal of winning it. Again.

When Aurora resolves to do ten push-ups a day and only does seven, this is non-ideal.

When Beth breaks her promise to attend her best friend's wedding simply to finish a book she's reading, this is non-ideal.

Many choices about food—what to eat, what to buy—are clearly not ideal. In Cormac McCarthy's *The Road*, some of the few human survivors of the apocalypse form cannibalistic gangs, roving the highways, capturing and eating other survivors. This is non-ideal because obviously morally wrong.¹

At least one of us invariably eats too much Halloween candy and ends up nauseated. This is non-ideal because imprudent.

One day after starting an attempt at vegetarianism, one of us had an Egg McMuffin, not realizing McMuffins contain ham. This is non-ideal because the policy of being vegetarian and eating ham is obviously incoherent.

Some food choices are non-ideal in, we think, more interesting ways. This paper is about some such departures from the ideal. Not all of these departures are *failures*. One might fall short of an ideal without failing. Some ideals are so difficult to achieve that simply approximating them counts as success: The ideal in baseball is getting on base every time, *never* making an out. So someone who only gets on base half the time is falling well short of the ideal; yet such a hitter is *extraordinarily* successful. When a doctor performs a less risky surgery at the cost of producing a less satisfactory result, this is

¹ A short account of human cannibalism is in Chapter Two of Fernández-Armesto (2002). A short philosophical discussion is in Diamond (1978).

non-ideal because it isn't the best possible outcome. But, still, the doctor mightn't have failed at all. She might have done just as she should have.

Also, not all failings to act ideally are fallings short. One might miss the ideal mark by exceeding the ideal. Eating a seventh burrito isn't ideal but also doesn't exactly fall short of some nutritional mark. Donating so much of your income to famine relief that you don't have enough left to feed yourself isn't ideal, but also doesn't fall short of the moral mark. (Though perhaps both of these are properly characterized as cases in which exceeding one standard - for example charity - results in falling short on another - for example, a duty of self-care.)

Our main goal in what follows is to map out (part of) the landscape - to identify, and provide examples of, several of the interestingly different ways in which we miss the ideal mark, food-wise. Some of these call for correction; some don't. Some are easy to fix; some aren't. Some are easy to see; some aren't. This paper isn't about fixing.² It is about how we get things not quite right: *believing* more or less than our evidence warrants, *under-* and *overcommitting*, and *under-* and *overperforming*. It is clearly, saddeningly, possible to fail in more than one of these ways at the same time - one can believe less than one's evidence supports, undercommit relative to one's moral beliefs, and underperform relative to one's moral commitments. By getting clearer on various ways we get things not quite right food-wise, we can figure out ways to improve, just as by getting clearer on whether it's the carburetor or the gas pedal or the windshield wipers that have gone wrong, we can figure out how to fix the car. But there's another lesson: some ways of not acting ideally are actually helpful and defensible. There's sometimes quite a bit to be said for getting things not-quite-right.

Carl knows he needs to get out of bed by 7. The ideal thing to do is sleep till then and bolt up. But, he knows, that almost never works for him, he finds getting out of bed too hard. He reluctantly sets his alarm for 6:45 to make sure he's out of bed by 7.

Delia knows the best thing for her to do is eat only a few pieces of Halloween candy. She also knows that once she starts, she usually can't stop and ends up with a bad stomachache. So she foregoes the best thing just to avoid the stomachache.

² For help with that, see Halteman and Zwart in this volume.

Eva thinks she should exercise more to be healthier. She treats herself to unhealthy candy every time she exercises. It would be better to just exercise, but Eva knows that that'd never work, that the unhealthy enticement is crucial to getting to the gym.

These behaviors are all non-ideal and all defensible, even laudable. Understanding the ways in which, and the reasons why, people can fall short of various ideals, can help to take a bit of the heat out of at least some of the debates about the ethics of eating. There's a sort of fury and superiority that attaches to some of the vegan/vegetarian rhetoric, which doesn't typically seem to help advance anybody's understanding of the issues, or improve anybody's behavior. (Often it hinders progress by making those in the crosshairs defensive and unreceptive to the substance of the criticism behind the outrage.) Sometimes this sort of harsh judgment of those in other camps might be justified, but other times not. Some people who aren't, for example, vegan might have excellent reasons for not being so. Relatedly, there's a sort of complacency and smugness about some omnivore rhetoric, often accompanied by a not especially helpful sort of outrage and/or contemptuous dismissal. Sometimes this sort of strong reaction might be justified, but other times not. The paper helps to pinpoint when and why.

1. DOING THINGS (MORE OR LESS IDEALLY)

Some people are dietary vegans, people who eat neither animals nor animal products. Some of these people never consider the issue and just eat that way because they were raised that way. Others carefully consider the issue and come to believe they *should not* eat any animals or animal products because it would be *morally wrong* to do so. So they commit not to do any such eating; they adopt a policy of doing so. And then they hold to that commitment, stick to that policy.

Other people are omnivores, people who eat animals and animal products as well as plants. Some of these people never consider the issue and just eat any old thing. Others carefully consider the issue and come to believe they *may* eat any old thing—more or less—because it is morally permissible to do so. So they adopt a policy of eating any old yummy thing, flesh or not. And then they hold to that policy.

What is going on here comes in three parts: There's what the agent believes he should or may do. There's a commitment to action, or adoption of a policy, based on

that belief. Finally, there's the action he then carries out based on that policy.

Not everything we do is so thoughtful or systematic. Some action is not produced by beliefs about what you should do or may do. When we idly drum our fingers on the table, this isn't produced by such beliefs.

Some action issues from a belief you should do it without issuing from a commitment. It's a one-off.

Other actions are repeated without being commitment-based. Where do you put your right ring finger when you are plugging in a microwave? For some people—for the two of us, at least—this action, though repeated, does not issue from commitments we make, or general policies we adopt, about how to act. Neither does our habitual behavior of checking the mailbox on the way into the office: It's something we do. We know we're doing it. We think it's the thing to do. But it doesn't issue from a belief about what we should do or a commitment to do such a thing—it's an act but also just a habit we have.³

Some things our bodies do aren't acts at all: hiccups, muscle twitches, fainting spells. These are things that happen rather than things we do.

Yet, we think, some of what we do with regard to food has the tripartite structure we introduced at the beginning of this section. We need to eat to survive. Sometimes, satisfying this need, we just eat. Other times, it's more complicated. We need to eat and then think about what to it: Should it be this or that? That or this? We consider the evidence. We form a belief about what we ought to eat. We set a policy: Veganism or omnivorism or the paleo diet. For some of us, these policies are an important part of who we are. For some of us, the policies match up with our beliefs about what we ought to do. For some of us, the

³ Here's a failure to act ideally that we will not be talking about: Sometimes you *should* form some commitment to act but do not do so. In such a case, you might act and this action isn't issuing from any sort of commitment, but, really, it should issue from a commitment. For example, you know you should limit your candy intake and know you should form a commitment to do this, a commitment to some sort of candy diet. But you don't. That's non-ideal. Similarly, you might fail to form a belief when, really, you should. You might, for example, be studiously uncommitted about whether the Earth's climate is changing despite having all the evidence that any reader of this book has. That's non-ideal. Such failures to attend to evidence and failures to form thoughtful commitments are important in food ethics. For discussion, see Nolfi (2015). What we focus on are mostly cases where you *do* believe, *do* commit, *do* act and these goings-on are non-ideal.

beliefs about what we ought to do match up with our evidence. The rest of the paper is about the rest of us.

2. BELIEF

Some deviations from the ideal come at the first stage when we form bad beliefs. A belief can be bad in all sorts of ways. We'll focus on beliefs that are bad because they are beliefs that are supported by evidence that warrants a *stronger* belief. For example, we learn that pandas are a type of mammal and that some people love all mammals. The evidence warrants the conclusion that some people love pandas. If, instead, we conclude only that some people love *female* pandas, we, to use an unlovely word, underbelieve.

For another example, we learn that the power is out. On this evidence, we are justified in concluding that we can turn on neither the TV nor the radio nor... If, instead, we conclude only that we can't turn on the TV, that's underbelieving.⁴

There's underbelief and then there's underbelief. When you look out your window and see a car, you might well have evidence about what its hubcaps are like, but, typically, you form no such belief. You thereby underbelieve. And that's okay. What the hubcaps are like is neither salient nor important in this case and so underbelief here isn't, we think, objectionable. It's non-ideal because it falls short of what an ideal cognizer would believe based on the same evidence, but falling short of this ideal is not, in this case, objectionable. Yet it is in some cases. Consider Fred. He believes it is morally wrong to buy factory-farmed meat. His argument for that conclusion is that factory farming produces a huge amount of suffering and what people get out of that sort of farming—food—is not sufficient justification for the suffering. Because of this, he

⁴ The philosophical terminology here might be confusing: The belief that some people love all pandas is stronger than the belief that some people love female pandas in exactly this way: That some people love all pandas implies that some people love female pandas. That some people love female pandas does *not* imply that some people love all pandas.

The belief that we can turn on neither the TV nor the radio is stronger than the belief that we can't turn on the radio in exactly this way: That we can turn on neither the TV nor the radio implies that we can't turn on the radio. That we can't turn on the radio does *not* imply that we can turn on neither the radio nor the TV. (Maybe the radio is broken and so can't be turned on but the TV works fine and can be turned on.)

thinks, *buying* the meat produced by factory farms is wrong.⁵ Basically, Fred reasons like this:

Factory farming causes a great deal of suffering for animals for very little gain for humans and animals. Hence,
It is morally wrong to factory farm. Hence,
It is morally wrong to buy factory-farmed meat.

Yet Fred's evidence for the claim that it is wrong to buy factory-farmed meat is evidence for the stronger conclusion that it is wrong to buy all factory-farmed food, not only meat but factory-farmed dairy and eggs, too. He believes morality requires him to X but does so on the basis of arguments that support a more-demanding requirement to X+. His evidence *obviously* supports the requirement to X+, and the requirement to X+ is a conclusion that's importantly relevant to what he's trying to figure out. In particular, Fred holds it is wrong to hurt animals just to produce cheap, yummy, nutritious food. Because of this, he holds it is wrong to buy animal *flesh* produced in that way. But if this argument that it is wrong to buy factory-farmed flesh is a good one, then it shows that it is wrong to buy any products of factory-farmed animals - it's not just meat that's problematic.

After all, meat animals aren't the only animals that suffer in factory farms in order to produce food for humans. Most obviously, there are hundreds of millions of factory-farmed egg-laying hens, hens who live in no better (and often worse) conditions than chickens raised for meat (and who live in such conditions for longer than meat birds do).⁶ There are millions of factory-farmed dairy cows, cows who live in no better (and often worse) conditions than

⁵ There are some suppressed premises in Fred's reasoning. One is the moral principle that licenses the move from the fact that factory farming causes a lot of suffering for very little gain to the conclusion that it's morally wrong to factory farm. The other is the principle that licenses moving from the wrongness of factory farming to the wrongness of buying factory farmed meat. Both principles are potentially contestable, but the second more so than the first. Fred's view is that *buying* factory-farmed products is typically wrong and his argument for this view flows from a premise about *factory farming* itself being wrong. Why is that a valid argumentative move? Why is it wrong to buy factory-farmed products just because the farming itself is wrong? This question is worth considering—and is considered in this volume in Terence Cuneo's, Elizabeth Harman's, Adrienne Martin's, and Ted Warfield's chapters.

⁶ The awfulness of certain forms of egg production is well-detailed in Singer and Mason (2007) and the documentary *Fowl Play*.

cattle raised for beef (and how live in such conditions longer than beef cattle do). These animals are hurt just in order to produce food. Unlike, say, the cattle that made your burger, they produce food without being killed, but Fred's argument for his food policy has nothing in particular to do with killing animals. Fred thinks it is wrong to cause animals much suffering in order to produce food. Because of this, he concludes it is wrong to buy factory-farmed meat. But he does not draw the more general conclusion that his premises warrant: That it's wrong to buy *any* factory-farmed food products. It is as if he is convinced by an argument the key premise of which is that causing pain to innocent people is wrong, and then, because he is thereby convinced, draws the conclusion that poking innocent people in the eyes is wrong, but then does not accept that kicking innocent people in the shins is wrong, too. Fred's is a common sort of failure, one of moral underbelief. Again, schematically, Fred believes that morality requires X, but on the basis of arguments that saliently support a more-demanding requirement to do X+.

A related failure is the failure of moral *overbelieving*. Consider Gabby. Gabby believes that eating meat is wrong on the basis of a belief that factory farming is wrong. If her argument is this simple

Factory farming is wrong. Hence,
All meat-eating is wrong,

she clearly draws a conclusion that goes well beyond what is warranted. Conclusions about *eating* don't follow straight from premises about *production*. You need a premise linking the wrongness of producing food a certain way to the wrongness of consuming such food. But even the more subtle

Factory farming is wrong.
Eating products that are wrongly produced is wrong. Hence,
All meat-eating is wrong

is not a good argument, and, if Gabby goes in for it, she believes more than her premises warrant. If factory farming were the only way of producing meat, then the argument would be better, but that isn't the only way. There are freerange-farmed animals. There are hunted animals. There are roadkilled-animals. Consider Jonathan McGowan, an English taxidermist. McGowan eats meat but only animals

that have been accidentally killed. Is he doing wrong? Gabby's premises don't show this. McGowan gets no meat from animal farming or hunting. (He, in fact, thinks both activities are morally wrong.) All he does is clean carcasses off roads and, rather than trashing them, consumes them.⁷

Because she overlooks the possibility of eating meat without eating meat that is wrongly produced, Gabby overbelieves and thereby fails to draw the food-based conclusion that her argument warrants. At most, from this argument, Gabby should believe that it is wrong to eat meat from places that wrongly produce it and that, in particular, it is wrong to eat factory-farmed meat. But that is consistent with it being permissible to eat meat produced in some other way.

3. COMMITMENT, PART ONE

Besides under- and overbelief, there's also the phenomenon of *under- and overcommitment*.

Consider Hope. She believes morality requires not eating any factory-farmed animals and because of this believes the ideal food policy for her is the policy of never eating a factory-farmed animal. Yet Hope *undercommits* by committing to a less-demanding policy of not eating any factory-farmed animals except those prepared at great restaurants and at family gatherings. If the steak tartare at Very Wonderful Restaurant comes from a factory-farmed cow, so be it, she'll eat it. If Grandma's turkey comes from a factory farm, so be it, she'll eat it.⁸

In one way, what's going on with Hope looks like what is going on with Fred; Hope seems to endorse an argument that looks like this:

Factory farming causes a great deal of suffering for animals for very little gain for humans and animals. Hence,
It is morally wrong to factory farm. Hence,

⁷ For more on McGowan, see <http://www.theguardian.com/lifeandstyle/2011/mar/26/i-eat-roadkill> and

http://www.huffingtonpost.com/2011/10/17/jonathan-mcgowan-roadkill_n_1016108.html

⁸ Almost all the turkeys do: 97.43% according to Foer (2009): 271. That's a good book for Hope since it is insightful on food and family.

It is morally wrong to eat factory-farmed products.⁹

So it might look like Hope makes a mistake like Fred's. It might look like she simply overlooks the fact that this argument implies that it is wrong to eat factory-farmed products at great restaurants and family gatherings. But this is an illusion. Hope makes no oversight. She knows full well that animals factory-farmed for great restaurants and for her grandmother's turkey casserole are no different from other animals. She knows full well that the conclusion of her argument is that eating factory-farmed food is wrong and knows full well that certain things she does are, by her own lights, morally wrong. Still she does them. And she doesn't do them absent-mindedly or out of weakness. Her failure is not one of performance or belief. She deliberately and consciously adopts a policy that she thinks is not the one that's best supported by the moral arguments. She thinks there's a sound moral argument against eating factory-farmed food, but she commits to a policy that permits eating factory-farmed food.

Because of this, the policy Hope follows is clearly, by Hope's own lights, morally inferior to a policy it is *possible* for Hope to follow. That is, there is an eating policy Hope could be pursuing—strict abstention from factory-farmed food—that is more coherent and, by her own lights, morally superior to the policy she in fact pursues. Hope *could* do this. She needn't eat any old meat at great restaurants or family gatherings. Abstaining would, however, be hard for Hope. It would require abstaining from her grandma's factory-farmed turkey casserole and foregoing the delicious factory-farmed steak tartare at Very Wonderful Restaurant.¹⁰ And so she commits to the less-difficult policy.

⁹ Maybe we need to make these arguments subtler in various ways in order to make it reasonable for Hope to be convinced by them. At any rate, keeping the arguments simple makes it easier to make our main points here, and we could make those points using more subtle arguments.

¹⁰ We assume that what we are *going* to do doesn't affect what we are *morally permitted* to do in the way that what we *can* do, arguably, does. So the fact that one of us is not *in fact* going to refrain from stealing this bar of soap does not show that it is morally permitted for us to steal the soap or even that it is *not wrong* to steal the soap. It is wrong. One of us is going to do wrong.

By contrast, the fact that one of us *cannot*—truly cannot—refrain from stealing the soap might show that it isn't wrong for that person to steal it. It wouldn't show that it is *permissible* to steal the soap; rather, it'd show that stealing isn't a morally evaluable action at

Doing so is a deviation from the ideal. For Hope commits to a policy that is, by her own lights, morally wrong and less coherent than a policy she could commit to. It could be her commitment is contemptible. It could be that what is going on with Hope is that she is "self-indulgent, essentially wishy-washy,...morally lazy...[and] suffers from the vices of ethical lethargy, hedonistic selfishness, and intemperance."¹¹ In some cases like Hope's, the grounds for the undercommitment are simply that committing to the more coherent, more morally defensible policy is difficult, and there is little to be said for such policies. For example, at least one of us promised himself to eat only 25 Christmas cookies last year. Could he have committed to doing better? Of course! But doing so would have been hard. So one of us undercommitted. That's not an impressive justification. When a young child tries to get out of tying his shoes on the grounds that doing so is hard, one might empathize without being impressed. Yet Hope's justification might well be quite a bit better than that, even quite good. It could be that her undercommitment, while non-ideal, is really quite defensible.

Why would anybody adopt a policy that, by their own lights, fell short of the requirements of morality? One reason: They just don't care about morality. Or - more likely - they just don't care *that much* about morality. More generously, they don't *just* care about morality, and they adopt a policy whose dictates are sensitive both to their moral views and to the other things they care about.¹⁴

all. It's like a reflex, a baby's instinctive action, or a meteor shower: something that happens that is neither right nor wrong.

See footnote fifteen for a bit more discussion and some references. For an introduction to issues about the relation between what you morally ought to do and what you can do, an argument about that relation, and a wealth of references to work on the topic, see Vranas (2007).

¹¹ This was an anonymous referee's description of her.

¹⁴ Generally, different values can pull you in different directions. You value community and self-interest. Should you join your friends in eating a tub of ice cream? There's a community-based reason to do so. There's a prudential reason not to do so. Should you keep serving your guest drinks? There's an etiquette-based reason to do so. There's a moral reason not to do so. (The example is from Foot (2002).) Sometimes it is easy to weigh up these competing values: Should you steal the apple just to see what it feels like? There's a prudential reason to do so. There's a moral reason not to do so. But the moral reason is obviously stronger—you shouldn't steal the apple. In other cases, the weighing up is tougher. In some cases, you might worry the values cannot be weighed against each other. On this, see Anderson

For example, a kid might adopt a policy of stealing cookies every once and a while even though she knows she shouldn't and might adopt it just because despite caring about not stealing and caring about doing right, she cares more about satisfying her own interest in having a cookie.

If Hope's motivation is like that, her actions are understandable but somewhat hard to defend. Her policy might, however, be better-defended. Her case might be like this: She thinks that following a morally sub-optimal policy is, morally speaking, the best that she probably *would* do. As we said, while the policy Hope follows is morally inferior to policies it is *possible* for her to follow, it is morally superior to the policies she would *actually* follow if she didn't follow hers. Hope, reflective, self-knowing, knows that if she failed to give herself the great restaurant escape clause and the family gathering escape clause, she would as a matter of fact follow policies she thinks are morally *much* worse: She'd too often give in to temptation, become discouraged, give up any effort at imposing moral constraints on her eating habits, and eat a totally unrestricted diet, including a lot of factory-farmed animal products. She is not, we stipulate, just eagerly seizing on any justification for not doing the more demanding thing. She is thoughtful, earnest, and undeceived. She has been down these roads before. She has powerful evidence about her track record with these sorts of commitments, and no evidence that this time will be different. Her confidence in what she would do is high and is warranted.¹⁵

So the policy she does follow—the one with the escape clauses—might be the best she *would* do – or the best she would be likely to do – out of the options that are

(1993) and Chang (1997) and the references therein. Also, see Ruth Chang's TED talk: https://www.ted.com/speakers/ruth_chang

¹⁵ Hope is in some ways like the sociopathic serial killer, Dexter Morgan, from the TV series, *Dexter*. Dexter follows a policy of killing only serial killers and other especially dangerous individuals. Part of his justification for this is that if he failed to kill serial killers regularly, he would do something by his own lights much worse: He would kill innocent people. (Hope, less so than Dexter, would benefit from reading the Halteman and Zwart essay in this volume.)

This example shows that a policy though morally *superior* to some alternative need not be morally *good*. This should be uncontroversial. Some actions are morally superior to others without being good. Stealing is morally superior to stealing and then lying about it, but stealing isn't morally good.

On the relations between what you *probably would* do and what you *should* do, see, among others, Goldman (1976), Goldman (1978), Jackson and Pargetter (1986), and Louise (2009).

realistic for her. That is, it might be the policy that is closest to the ideal, the one with the fewest and least serious deviations. Yet following that policy isn't the best she *could* do. Again, it is *possible* for her to strictly abstain from factory-farmed food. Just like it'd be possible for the two of us to do 100 push-ups a day or run marathons every year. It's possible. It's just not going to happen. And it's not crazy to think that it's reasonable for us to make our plans and decisions, and choose our policies, based on the assumption that it's not going to happen. Perhaps something stronger is true: It would be unreasonable for us *not* to make our plans and decisions, and choose our policies, based on the assumption that it's not going to happen.

Here's a way of making the point: Suppose we have the option of signing up to serve on an alpine rescue team next summer. It's a high-stakes job, and one that requires a very, very high level of fitness. We could make it the case that we're at that level of fitness by the time next summer rolls around, by adopting and sticking to a very, very demanding training regimen. We could do it! Yet we both know that neither of us will in fact stick to such a regimen, and so we won't have the level of fitness that we'd need to have in order not to be dangerous liabilities on the alpine rescue team. The *ideal* thing for us to do is to sign up for the job, do the fitness program, and then go do a great job as alpine rescuers. And it's *possible* for us to do that. But we are quite reasonably very confident that that is not what will actually happen if we sign up for the job. If we sign up for the job, we will fail to stick to the fitness program, and we'll be dangerous liabilities as alpine "rescuers".

The alpine rescue case is similar to (because based on) Jackson and Pargetter (1986)'s Professor Procrastinate case, a case in which a professor is asked to review a book, knows that the best thing for him to do is to say "yes" and then write the review, and that the worst thing is to say "yes" and not write the review, and confidently predicts that if he says "yes", he won't write the review. The Procrastinate case, the mountain rescue case, and the class of cases of choosing a food policy that we're presently concerned with all have the same troublemaking structure: There's a commitment one can choose to undertake now, or not, and there's a following-up on the commitment that one can perform later, or not. The best thing to do is to undertake the commitment and then follow up. The worst thing to do is to undertake the commitment and then fail to

follow up. That precipitates the review not being written and the editor and author being left in the lurch. It precipitates incompetent and dangerous mountain rescue workers endangering people's lives. It precipitates the agent with Hope's sort of motivational structure becoming discouraged and giving up on the project of improving moral dietary improvement altogether. The case for choosing the second-best, and undercommitting relative to what you take to be the morally best course, is that you expect yourself to do better, morally speaking, by adopting the second-best commitment than by trying, and failing, to abide by the best commitment. The best you will do is second-best.¹⁷

Our knowledge that we won't, in fact, follow through on the optimal course can quite appropriately inform our decisions about what to do, and so that the thing for us to do is to choose the second-best option of not signing up for the alpine rescue team in the first place. Something like this idea is at work when G.E.M. Anscombe claims that some people think that they should be pacifists, try to be pacifists, find that they are not living up to their ideals and, disheartened, give up on even approximating the ideals. Even if one acknowledges that pacifism is the ideal, committing to a less pacific life might make sense.¹⁹

Peter Singer, in his work on famine relief, makes a similar point. His view is that the affluent should give up quite a bit to help the poor. How much? He doesn't make too many demands of the affluent for he fears that they will see what his view requires, come to believe it's right, and, since they aren't going to comply with it, just say "forget it" and ignore the poor entirely.²⁰

¹⁷ Some moral theories imply that each of us is morally required to do what is best. Others are skeptical that *the best* makes sense. See Thomson (1997) and (2008).

Still others accept that some actions are best but imply that some actions are morally permissible—or even required—even though they are not best. These theories sometimes carve out space for a category of action that is *supererogatory*, permissible, praiseworthy, but above and beyond the call of duty. For an overview with lots of references, see Heyd (2012)

¹⁹ Anscombe writes, "The truth about Christianity is that it is a severe and practicable religion, not a beautifully ideal but impracticable one." The false view of Christianity moves pacifists and "pacifism has corrupted enormous numbers of people who will not act according to its tenets. They become convinced that a number of things are wicked which are not; hence seeing no way of avoiding wickedness, they set no limits to it" ((1961):56-57).

²⁰ See Singer (1972) and (2010). In the latter, Singer writes, Asking people to give more than almost anyone else gives risks turning them off, and at some level might cause them

Abelard suggests that a policy of assiduously avoiding minor sins leads to a commitment of major sins. One plausible way of explaining why: Since, realistically, you're going to commit a minor sin every now and again, straining to avoid them might just weaken your self-control so much that you end up committing a major one.²¹

One of the ideas behind the claim that the perfect is the enemy of the good is that striving for an ideal can be counterproductive in that it can stop you from doing even pretty well and lead to your, instead, doing worse. Hope's food policy is pretty good, though imperfect.²²

4. COMMITMENT, PART TWO

Choosing the not-quite-best is a way of undercommitting oneself. It's also possible - and it's plausibly sometimes *reasonable* - to overcommit yourself. You might believe morality requires merely X and nevertheless adopt a more-demanding policy of doing X+.

Parents commonly adopt more restrictive policies than they think they need to. They are opposed just to letting their youngest kids watch certain R-rated movies but impose a policy of no watching R-rated movies. They are opposed just to bags of Cheetos-consumption but impose a policy of no Cheetos-consumption.

to question the point of striving to live an ethical life at all. Daunted by what it takes to do the right thing, they may ask themselves why they are bothering to try. To avoid that danger, we should advocate a level of giving that will lead to a positive response. (*Ibid.*: 151)

²¹ Abelard (1995): 29-33. Thanks to William Mann for the reference. Note that a policy of *routinely* going in for minor sins might also weaken one's self-control and one's resolve to resist major sins. On this, see the *Catechism of the Catholic Church* 1876: "The repetition of sins - even venial ones - engenders vices, among which are the capital sins." This sort of issue is discussed in the next section.

²² As we said, there are ways of describing the Hope case so that her policy is not so good. It's an important part of the case that she knows that following this policy is the best she's likely to do. What if we change that detail? Imagine if Hope simply deceives herself into thinking that the policy she does follow—and is very content to follow!—is the best she *would* do when, in fact, she could and would (if she set out to) follow a more severe, less pleasurable policy. She might, for example, tell herself that in all the nearby possible worlds where she tries to be an exceptionless freer vegetarian, she fails and ends up stuffing her face with McNuggets. But she might be fooling herself about that or simply wrong: There might be some nearby worlds in which she's a strict freer vegetarian. But these worlds are less appealing to her since they involve no turkey at Thanksgiving, no steak at Very Wonderful Restarant. That saddens her. And that affects her view of what would happen. If that's what's going on, then Hope's policy becomes no less understandable but quite a bit less impressive.

The monks of Mount Athos generally eat a significantly more restrictive diet than they think their faith requires of them.²⁴

And consider Isaiah. He believes it is typically wrong to kill sentient animals. Because of this, he believes it is typically wrong to eat sentient animals.²⁵ However, the policy he follows is not of abstaining merely from eating *sentient* animals killed for food. Rather, the policy he follows is not eating any animals: He doesn't eat non-sentient animals that have been raised and killed for food, things like farmed clams. Isaiah also doesn't eat sentient animals that weren't killed for food, such as road-killed deer or wild boar, struck by lightning.

Like Hope's, Isaiah's case involves a mismatch between moral argument and behavioral policy. Unlike Hope's, Isaiah's policy is *more* restrictive than the policy that his argument supports. All the same, there is something non-ideal about it: By not eating various foods that, by his own lights, it is morally permissible to eat, Isaiah deprives himself of lots of pleasurable experiences that are not, by his own lights, wrongful. He also imposes costs on others: He is unable to coordinate with his fellow diners on a shared order of mussels to accompany dinner, for example, and it's not enough, for a restaurant to qualify as a place to meet Isaiah for lunch, that it does a great oyster po-boy. That restaurant has to offer, for example, a really good seitan burrito. Isaiah's argument permits him to eat more (kinds of) food than he in fact does. Isaiah is, like Hope, adopting a policy that he doesn't regard as lining up with the moral facts.

Yet, like Hope's, Isaiah's position, while not by Isaiah's own lights ideal, is defensible. In fact, like

²⁴ When they fast, the monks

...generally follow a strict pattern and abstain from any consumption of meat (which monks never eat anyway), dairy products, fish, wine, and oil. The notion is that through fasting the body shares in the work of prayer; and fasting gives the faithful a sense of freedom and lightness... The monastic diet is much the same as the traditional Greek peasant diet and consists largely of vegetables, bread, olives, salads, cheese, soya and fruit. On non-fast days there is usually plenty of food on the monks' table and an ample supply of wine; and on feast days there is often fish. (Speake (2002): 222)

²⁵ For that matter, he doesn't eat sentient animals that were killed for certain other reasons—for fun, say—and then offered up as food.

Hope's, Isaiah's policy makes a great deal of sense.²⁶ How so?

Generally, adopting an over-demanding policy makes sense if you predict that you will underperform if you adopt the less demanding, morally ideal policy and predict you'll be able to do better if you adopt the more demanding one. Maybe you think you'll do this because you think you're just the kind of person who systematically underperforms relative to commitment, and so you need to overcommit in order to perform on target. Maybe you think this because you think underperformance (or particular kinds of underperformance) is especially likely in this case. For example, you think it's okay to eat humanely raised meat a couple of times a week. But if you adopt a policy of doing that, you'll eat humanely raised meat five or six times, and a factory-farmed burger once or twice. If you adopt a vegetarian policy, you'll still fall short, but the way you'll do that is by eating freerange meat once or twice a week. For a non-food-related example, you might commit to going to the gym ten times a month because you're happy with going eight times but know that you also won't quite meet whatever goal you set yourself.²⁷

It's fairly common to find vegans or vegetarians who don't think that absolutely all meat/animal product-eating is impermissible. They think there is nothing wrong with eating roadkill or eating discarded food the eating of which definitely won't contribute to demand.²⁸ Some even

²⁶ The policy might be more defensible than Hope's. It might be an Anscombean rejoinder to Hope, a "severe and practicable" policy that reveals that she lets herself off the moral hook too easily. Whether this is so depends on whether Hope is right about what she would do if she weren't a freerange vegetarian with exceptions. Would she actually eat any old factory-farmed thing? Or could she stick to strict freerange vegetarianism?

²⁷ By adopting more demanding policy, you guard against a more insidious sort of underperformance. This comes from the phenomenon Richard Holton calls 'judgment shift' in which what temptation does is not just get you to want something but also to adjust your valuation of whether you *should* do it. So, for example, when you're tempted by cigarettes, sometimes what happens is you know you shouldn't smoke but just can't help yourself, then smoke. (We discuss cases like this in §5.) Yet temptation sometimes works differently: Sometimes what happens is that the temptation gets you to revise your view about what you should do, so you end up thinking, "It's fine for me to smoke this one." So you do. This is underperformance by the lights of your past self rather than underperformance by lights of your current self. See Holton (2009) on judgment shift and Korsgaard (1989) and Parfit (1984) on present underperformance by lights of your past self.

²⁸ See Rachels (2011) and some papers discussed in Warfield's paper in this volume.

think there is nothing wrong with eating the products of certain presently-uncommon methods of animal agriculture. But applying the policy that tracks this tortuous moral boundary is cognitively demanding, requires more research, introduces space for self-deception and rationalization, and introduces opportunities for bad action based on bad information or misapplication of the more complex rule. By adopting a strict, say, vegetarian policy, one avoids a lot of complication in one's food decisions.

For the over-demanding policy might be easier to *stick to*: The clear, never-to-be-crossed bright line makes the policy easy to follow and, importantly, easier to follow than the ideal policy, a policy that casts a line of less brightness.

The more-demanding policy might be psychologically easier to *keep in mind*: It's not just a bright line. The bright line is also cast by a simple, easy-to-keep-in-mind rule.

The more-demanding policy introduces fewer potential sources of error: If it's an animal, you can't eat it. If it isn't, you can. By contrast, the less-demanding policy leaves you open for making mistakes on the basis of all sorts of errors about the provenance of this particular bit of flesh: Was the animal sentient? Was it killed for food? Etc. And avoiding those errors requires research. By contrast, the more-demanding policy necessitates no check on the provenance of what you're eating—whether the animal was killed for food, for example. Also, there's no need to check on the mental status and cognitive capacities of the animal—whether, for example, it's sentient, like a cow, or not, like a mussel. And no need to put trust in research that is not settled.²⁹

Adopting the more-demanding policy might help in avoiding temptation. Attempting to adhere to the ideal policy might tempt Isaiah to act badly. (Isaiah's more restrictive policy puts up a buffer between him and the animals he thinks it's actually wrong to eat.)³⁰ If Isaiah allows himself to eat any sentient animal so long as it isn't killed for food, he might get into the habit of eating, say, road-killed chicken. This might tempt him to

²⁹ Allen and Trestman (2014) covers recent research on animal cognition. It gives a good feel for how interesting and tricky it is and suggests that when we make mistakes about which things are sentient, these mistakes are typically cases of false negatives: We falsely believe certain things aren't sentient and they turn out to be sentient.

³⁰ Compare the Talmud's advice to "make a fence around the Torah" (Avot 1:1). Thanks to Arthur Kuflik for the reference.

eat non-road-killed chicken: If you're eating chicken already, why would this one farm-raised chicken hurt? So the ideal policy might be harder to adhere to because it makes certain deviations so tempting.

There is also likely to be an *epistemic* advantage to the more-demanding policy: The more demanding policy might well be one Isaiah is **certain** is okay. Even if he is pretty sure eating Niman Ranch³¹ bacon is permissible, he might still want to play it safe because he is not **completely** sure, and it's a big deal if he's wrong.³² Compare: You are somewhat sure it's safe to eat food weeks after the sell-by date but are much more sure it's safe to eat food by the sell-by date. It makes sense to adopt the more demanding policy of eating food by the sell-by date. (It's more demanding because you have less time to get your eating in.)

Also, in some cases following the more-demanding, more cautious, bright line policy will be likely to have better effects on others than the ideal policy. Peter Singer thinks some meat-eating is permissible but instead follows a vegan diet in part because that has better effects on others than a freerange one.³³ If you're visibly a vegan, this might lead to conversations with one's dining companions in which they might be exposed to ideas, arguments, and examples that change their behavior, conversations they wouldn't go in for if you were following the ideal eating policy.

Also, eating meat very rarely and under very strict conditions might signal to others that all meat-eating is okay, and that might be a bad effect. When Isaiah and Singer instead adopt a policy of never eating meat, they signal that they are against eating animals. In the US, at least, not eating meat signals that one is against killing animals. These are the signals that are in fact picked up whether or not they are the signal Isaiah and Singer mean to send and whether it is a signal they endorse. Sending out this signal might make sense even if a policy of eating animals sometimes has more to be said for it than a policy of never doing so. When you signal that, you signal you are on a certain side of an issue. There might be value in that that cannot be reduced to the value of its effects on

³¹ On Niman Ranch, see Foer (2009) and Niman (2010). For more on similar farms, see Kingsolver (2007).

³² Cf. Guerrero (2007).

³³ Cf. Singer (1999).

Isaiah or Singer or the effects on those who pick up the signal.³⁴

5. PERFORMANCE

Besides missing the mark with our belief and commitment, we can miss the mark with our performance..

We focus on *underperformance*.³⁵ Juan believes that morality requires that he do X, he commits to a policy of doing X, but he doesn't do X. There are more and less interesting ways this might happen. For example, Juan believes he shouldn't eat Cheetos, he commits to a policy of not eating Cheetos, and then, when he's out and about, he eats some trail mix with well-hidden Cheetos in it. This is underperformance due to (understandable) confusion. Not so puzzling.

In another not-so-puzzling example, Juan believes he shouldn't eat Cheetos, he commits to a policy of not eating Cheetos, and then, when he's out and about, someone puts a gun to his head and tells him that he had better eat Cheetos. So he does. This is underperformance due to coercion. Again, not so puzzling.

In a third example, Juan believes he shouldn't eat Cheetos, he commits to a policy of not eating Cheetos, and then, when he's out and about, he's given a drug that makes him irresistibly crave Cheetos. So he eats some. This is underperformance due to addling. Again, not so puzzling.

A more puzzling example of underperformance is this: Juan believes it's morally wrong to eat Cheetos, that he shouldn't—morally shouldn't—eat Cheetos, he commits to a policy of not eating Cheetos, and then, when he's out and about, under no coercion or illusion, not addled, harboring no relevant false beliefs about them, he eats a bunch of Cheetos.

What is going on here? What exactly is happening in situations where we do something other than the thing that we sincerely judge we should do? To answer this is to solve the traditional philosophical puzzle of *akrasia* or weakness

³⁴ On signaling and symbolism, see, among many others, Adams (2002) and Andrew Chignell's, Terence Cuneo's, and Adrienne Martin's papers in this volume.

³⁵ Compare with overperformance. Overperforming is believing you should X but need not X+, committing to X rather than X+, and then doing X+. You overperform when you believe you should clean your office once or twice a month, commit to doing so, and then clean it daily. ("I can't stop.") We focus on underperformance because, as far as we can tell, dietary overperformance is usually pretty *philosophically* uninteresting and unproblematic. Counterexamples are, of course, always welcome.

of will.³⁶ Juan's Cheeto-eating and other, less frivolous food cases raise a puzzle right next door to the traditional puzzle, one about cases in which we fail to follow through on our sincere commitments or intentions.

There are a number of different things that could be happening. Here is a (very incomplete!) sample:

Juan might think that he *morally* should not eat the Cheetos, but think that other, non-moral considerations outweigh the moral ones, so all-things-considered he should eat them. Reasoning like this goes on with regard to veganism sometimes. People are sometimes convinced that *morally* they should be vegan and also think that nutritionally veganism is inadequate for them and so think that veganism is *imprudent*. They then think the prudential considerations are weightier than the moral ones and decide that, all things considered, they should not be vegan. Such cases do not, it turns out, raise a puzzle about failing to carry out one's sincere commitments. In such a case, it turns out, one's sincere commitment is to eat the Cheetos or not to be vegan and that's what one does. But such cases do raise questions about how forceful moral considerations are. Sometimes those considerations seem to be decisive? "Should I lie to him?" "No. That would be wrong." But other times, maybe they aren't. When aren't they? Why?

Alternatively, Juan might be fragmented - part of him thinks he should X, part of him thinks he should Y, and it's the first part that adopts the policy and the second that pulls the behavioral trigger at the crucial moment. We aren't thinking of Juan as having a split personality. Rather, we think of him as being a normal, complex person who's torn about something. He's like you when you get home from some terrifying movie. You know there is no goblin in your basement. Yet you find yourself unable to go down to it to scoop the kitty litter. There's a rational part of your soul that's forming a belief. And an arational part of it that's driving your behavior, keeping you from going downstairs. Juan might show a different but related split. So, too, might people who keep botching their diets or, to return to an earlier example, keep eating too many cookies.³⁷

³⁶ Whether akrasia and weakness of will are the same thing is disputed. For discussion, see Holton (2009) and Stroud (2014).

³⁷ For this sort of story about akratic action, see Davidson (1970, 1982). For other discussions of the general phenomenon of fragmentation, see for example Lewis (1982), Stalnaker (1984), Schwitzgebel (2001), Egan (2008), Greco (forthcoming), and Elga and Rayo (ms).

Alternatively, Juan's behavior might be the product of a behavior-guidance system that's not sensitive, or not just sensitive, to Juan's reflectively endorsed ends, and his reflectively endorsed beliefs about what's likely to best satisfy those ends. J.M. Coetzee suggests that "the level of behavior [Juan wants] to change is too elementary, too elemental, to be reached by talk."³⁸ There is a fair bit of research suggesting that a lot of our eating behavior is driven by systems of behavior guidance that largely short-circuit our rational, reflective systems of deliberation and decision.³⁹ You might have a commitment to not eating any more hors d'oeuvres at the party, but your eating of them is insensitive to that, automatic.

Finally, Juan might be discovering that his commitment, after all, wasn't really a commitment not to eat Cheetos. It was, instead, some more subtle commitment to mostly avoid eating Cheetos or to avoid eating Cheetos when at home or... Relatedly, imagine you find yourself both avowing that you eat no animals and yet also occasionally eating oysters or mussels. This could be a case of not following through on your commitment or, instead, it could be a case of discovering that you have been misunderstanding or misdescribing your own commitments: What you're against is eating *sentient* animals rather than eating animals.

Plausibly, there's no one-size-fits-all account to be given of all cases in which we fail to do what we've committed to doing. Plausibly, different theories explain what is happening to different people on different occasions of moral underperformance. Plausibly, too, more than one theory might correctly characterize the same person on the same occasion. Because of this, the steps to take to deal with underperformance—your own or someone else's—will vary from case to case.

6. CONCLUSION

Unlike breathing or sleeping, eating is both an activity we need to do to survive and one that routinely raises pointed ethical questions. It might well be that the eating policies that dominate the globe are morally wrong and *ipso facto* non-ideal. It might be that the food production policies that dominate the globe are wrong and *ipso facto* non-ideal. If so, multiple times a day, huge numbers of people act non-ideally and objectionably so.

³⁸ Coetzee (2004): 103. The point recurs several times in the book.

³⁹ See Farley and Cohen (2009) and Wansink (2010) and the references therein.

Also, it might be that the right policy will be hard to follow, the wrong one will always tempt. Compare this with, say, stealing or killing innocent people. Doing so is non-ideal, but comparatively few people, comparatively infrequently, are tempted to steal or to kill people.⁴⁰ By contrast, eating meat might be non-ideal and powerfully tempting. Generally, acting non-ideally food-wise might be tempting in a way that acting non-ideally theft-wise or murder-wise is not.

In this paper, we have identified, and provided examples of, several of the interestingly different ways in which we can fall short of an ideal, be that ideal moral, prudential, or rational. We also provided ways in which we can fail to hit the mark by overshooting it. We can *underperform* relative to our commitments: We can identify some policy as best, set out to follow it, and then fail to do so. So, too, we can *overperform*: We can identify some policy as best and some more stringent policy as above and beyond the call of duty, we can set out to follow the less stringent policy and yet find ourselves following the more stringent policy. We can *undercommit* relative to our beliefs: We can identify some policy as best but then set out to follow some policy that is by our own lights less stringent. And so, too, we can *overcommit*: We can identify some policy as best but then set out to follow a policy that is by our own lights even more stringent. Finally, we can *underbelieve* relative to what the considerations we bring to bear on our thinking support: We can come to believe less than our evidence supports. And so, too, we can *overbelieve*: We can come to believe more than our evidence supports.

By identifying some different ways we can act in non-ideal ways, we are trying to identify different ways we go wrong and are tempted to go wrong. This taxonomical project helps to make sense of why people act in ways that we mightn't immediately understand and yet, on reflection, might make good sense.⁴¹ The taxonomical project helps, too, with fixing these deviations from the ideal when they crop up. It helps formulate more realistic policies, too,

⁴⁰ That said, the ideal of not killing humans may well have become easier and easier for people to follow, and maybe centuries from now people will look back at our food practices with the sort of horror and incredulity with which we now regard cultures in which routinely killing people is normal.

⁴¹ Compare this with Elizabeth Harman's project of making sense of a certain sort of vegetarian accommodation. See her paper in this volume and also her (2015b).

policies that are easier to follow and easier to stick to. These policies accommodate the fact that we are weak and limited in various ways. Because of this, there will be some deviations from the ideal. There are better and worse ways of deviating. *Worstward Ho's* advice to "fail better" is sound.⁴²

⁴² "All of old. Nothing else ever. Ever tried. Ever failed. No matter. Try again. Fail again. Fail better" (Beckett (1983): 7). This paper fails in fewer ways than it would have due to the efforts of friends who commented on drafts, suggested readings, and improved our ideas. Thanks, Andrew Chignell, Terence Cuneo, Matthew Halteman, Elizabeth Harman, Arthur Kuflik, and Kate Nolfi.

Bibliography

- Abelard, Peter. 1995. *Ethical Writings*. Indianapolis, IN: Hackett.
- Adams, Robert. 2002. *Finite and Infinite Goods*. New York: Oxford University Press.
- Allen, Colin and Trestman, Michael. 2014. "Animal Consciousness", *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Edward N. Zalta (ed.), forthcoming URL = <http://plato.stanford.edu/archives/sum2014/entries/consciousness-animal/>.
- Anderson, Elizabeth. 1993. *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press.
- Anscombe, G.E.M. 1991. *Ethics, Religion, and Politics*. Oxford: Blackwell.
- . 1961. "War and Murder" in Anscombe (1991).
- Beauchamp, Tom and Frey, R.G., ed.s. 2011. *The Oxford Handbook of Animal Ethics*. New York: Oxford University Press.
- Beckett, Samuel. 1983. *Worstward Ho*. London: John Calder.
- Chang, Ruth, ed. 1997a. *Incommensurability, Incomparability, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Chang, Ruth. 1997b, "Introduction", in Chang (1997a).
- Coetzee, J.M. 2004. *Elizabeth Costello*. New York: Vintage.
- Cuneo, Terence. 2015. "Conditional Moral Vegetarianism" in this volume.
- Davidson, Donald. 1970. "How is Weakness of the Will Possible?" in Feinberg (1970).
- Davidson, Donald. 1982. "Paradoxes of Irrationality" in Wollheim and Hopkins (1982).
- Diamond, Cora. 1978. "Eating Meat and Eating People," *Philosophy* 53: 465-479.
- Egan, Andy. 2008. "Seeing and Believing: Perception, Belief Formation and the Divided Mind," *Philosophical Studies* 140: 47-63.
- Elga, Adam and Rayo, Agustín. MS. "Fragmentation and Information Access"
- Farley, Deborah and Cohen, Thomas. 2008. "Eating as an Automatic Behavior," *Preventing Chronic Disease* 5(1): A23.
- Feinberg, Joel, ed.. 1970. *Moral Concepts*. New York: Oxford University Press.

Fernández-Armesto, Felipe. 2002. *Near One Thousand Tables*. New York: Free Press.

Foer, Jonathan Safran. 2009. *Eating Animals*. New York: Little, Brown.

Foot, Philippa. 2002. "Are Moral Considerations Overriding?" in *Virtues and Vices*. New York: Oxford University Press.

Goldman, Holly Smith. 1976. "Dated Rightness and Moral Imperfection," *Philosophical Review* 85: 449-487

--. 1978. "Doing the Best One Can" in A. Goldman & J. Kim (Ed.), *Values and morals*. Dordrecht: Reidel.

Greco, Daniel. Forthcoming. "Iteration and Fragmentation," *Philosophy and Phenomenological Research*.

Guerrero, Alexander. 2007. "Don't Know, Don't Kill," *Philosophical Studies* 136: 59-97.

Halteman, Matthew and Zwart, Megan. 2015. "Philosophy as Therapy for Recovering (Unrestrained) Omnivores" in this volume.

Harman, Elizabeth. 2015a. "Eating Meat as a Morally Permissible Moral Mistake" in this volume.

--. 2015b. "Morally Permissible Moral Mistakes," forthcoming in *Ethics*

Heyd, David. 2012. "Supererogation," *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2012/entries/supererogation/>.

Holton, Richard. 2009. *Willing, Wanting, Waiting*. New York: Oxford University Press.

Jackson, Frank and Pargetter, Robert. 1986. "Oughts, Options, and Actualism," *Philosophical Review* 95: 233-255.

Jamieson, Dale, ed. 1999. *Singer and His Critics*. Boston: Wiley-Blackwell.

Kingsolver, Barbara. 2007. *Animal, Vegetable, Miracle*. New York: Harper.

Korsgaard, Christine. 1989. "Personal Identity and the Unity of Agency: A Kantian Response to Parfit," *Philosophy and Public Affairs* 18: 101-132.

Lewis, David. 1982. "Logic for Equivocators," *Noûs* 16: 431-441.

Louise, Jennie. 2009. "I Won't Do It! Self-Prediction, Moral Obligation, and Moral Deliberation," *Philosophical Studies* 146: 327-348.

Martin, Adrienne. 2014. "Factory Farming and Consumer Complicity" in this volume.

Niman, Nicolette Hahn. 2010. *Righteous Porkchop*. New York: William Morrow.

Nolfi, Kate. 2015. "Food Choices and Ethical Character," manuscript

Parfit, Derek. 1984. *Reasons and Persons*. New York: Oxford University Press.

Schwitzgebel, Eric. 2001. "In-Between Believing," *Philosophical Quarterly* 51: 76-82.

Singer, Peter. 1972. "Famine, Affluence, and Morality," *Philosophy and Public Affairs* 1: 229-243.

--. 1999. "Reply to Hare" in Jamieson (1999).

--. 2009. *The Life You Can Save*. New York: Random House.

-- and Mason, Jim. 2006. *The Ethics of What We Eat*. Emmaus, PA: Rodale.

Rachels, Stuart. 2011. "Vegetarianism" in Beauchamp and Frey (2011).

Speake, Graham. 2002. *Mount Athos: Renewal in Paradise*. New Haven, CT: Yale University Press.

Stalnaker, Robert. 1984. *Inquiry*. Cambridge: MIT Press.

Stroud, Sarah. 2014. "Weakness of Will," *Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), URL = <http://plato.stanford.edu/archives/spr2014/entries/weakness-will/>.

Thomson, Judith Jarvis. 1997. "The Right and the Good," *The Journal of Philosophy* 94: 273-298.

--. 2008. *Normativity*. Chicago: Open Court

Vranas, P.M.B. 2007. "I Ought, Therefore I Can," *Philosophical Studies* 136: 167-216.

Wansink, Brian. 2010. *Mindless Eating*. New York: Bantam.

Warfield, Ted. 2014. "Eating Dead Animals: Meat Eating, Meat Purchasing, and Proving Too Much" in this volume.

Wollheim, Richard and Hopkins, James, ed.s. 1982. *Philosophical Essays on Freud*. Cambridge: Cambridge University Press.