

**A Case for Ethical Veganism:
Intuitive and methodological considerations.**

Tristram McPherson
Virginia Tech Department of Philosophy
tristram@vt.edu

Penultimate Manuscript for *Journal of Moral Philosophy*
When the paper is published, it will be available at:
<http://www.brill.com/journal-moral-philosophy>

Abstract

This paper begins by setting out an intuitive case for ethical veganism: the thesis that it is typically wrong to consume animal products, that begins with the intuitive claim that it is wrong to set fire to a cat. I then raise a methodological challenge: this is an intuitive argument for a revisionary conclusion. Even if we grant that we cannot both believe that it is permissible to drink milk, and that it is wrong to set fire to cats, this leaves open the question of which of these judgments we should abandon. I consider and reject three strategies for addressing this question: more methodologically naïve moral theorizing, appeal to systematic normative theory, and attacking non-moral presuppositions. I argue that philosophically satisfying resolution of the conflict requires debunking our grounds for belief in one of the conflicting claims. Finally, I argue that ethical veganism is supported by consideration of the most salient debunking arguments available.

Keywords: animal ethics, applied ethics, methodology in ethics, veganism.

Introduction

One of the central features of human life is that we interact with non-human animals. We observe them; we occasionally compete with them for resources; some of us form relationships with them as pets; and almost all of us eat, wear, and otherwise use products made from them or by them. This paper argues that the latter interactions are morally problematic: it is typically wrong to use or eat products made from or by certain familiar non-human mammals, including cats, dogs, cows, pigs, sheep, and deer.

I call this thesis *ethical veganism*. This thesis is weaker than standard forms of veganism in two respects. First, I argue that it is typically wrong to use animal products, but I do not think that this prohibition is exceptionless. Second, my argument rests in part on animals' capacity to suffer. In this paper, I do not address the difficult question of how far this capacity – and hence the scope of my argument – might be extended from the familiar mammals mentioned to various other species. However, it certainly does not apply to the entire animal kingdom, as some readings of veganism would require. For example, I am certain that it does not apply to bivalves, which lack a brain, or sponges, which entirely lack a nervous system.¹ For convenience, I will nonetheless use the word 'animals' to refer to the familiar mammals listed above.

The paper begins by sketching a direct intuitive argument for ethical veganism (§1-2). This thesis is *revisionary*: using animal products is a ubiquitous feature of human life, and one that appears morally innocuous to most of us. I show that this fact poses a neglected but important methodological objection to direct intuitive arguments for veganism: at best, such an argument can show that some of our intuitive judgments about animal ethics are inconsistent. This does not yet show that ethical veganism is the uniquely reasonable resolution of this inconsistency (§3). In §4, I consider strategies for resolving the intuitive inconsistency. I argue that part of the most philosophically satisfying way to do so appeals to asymmetries in the debunking arguments available to undercut the intuitive judgments that ground the opposing positions. In §5 I argue that ethical veganism is supported by the greater plausibility of the debunking argument available to it. I conclude that veganism is supported by the balance of evidence.

¹ For a case for using 'vegan' in something closer to the sense that I intend, see Cox 2010.

This paper is structured as an extended argument for that conclusion. However, it can also be read as a case study in the methodology of applied ethics. There are *prima facie* attractive intuitive arguments for many revisionary views in applied ethics. These arguments will potentially face methodological challenges parallel to the one that I pose to the vegan in §3. The details of how best to assess such challenges will vary in crucial detail from case to case. However, the discussion of §4 and §5 provides a framework and a model for addressing such cases.

1. The intuitive case against causing animal suffering

Some leading discussions of animal ethics implicitly or explicitly presuppose a systematic normative ethical framework.² The argument of this section and the next will instead deploy a familiar model of philosophical theorizing: what I will call the *intuitive-explanatory* model. On this model, one appeals to intuitively compelling judgments about clear cases, and seeks to construct local ethical principles capable of explaining the truth of those judgments, without appeal to systematic normative or metaethical theory, or more complex methodological strictures.³

This section argues that it is typically wrong to cause animal suffering. The argument begins with an intuitively compelling claim about a case:

Cat It is wrong to pour gasoline on a cat and light it on fire.⁴

Like most of the other ethical claims that I will be making in this paper, Cat should be read as including an implicit ‘other things being equal’ clause. Things would not be equal, for example, in the following variant on one of Bernard Williams’ famous examples (1973, 98-9): suppose a militia captain credibly threatens Jane that he will refrain from executing thirty innocent villagers just in case she sets his cat on fire. I take it that Jane should obviously set the cat on fire, and I do not intend Cat to be incompatible with this conclusion. To make the plausibility of Cat vivid, suppose that I came to a philosophy

² Explicit examples include Baxter 1974, Regan 1983, Korsgaard 2004, Wood 1998, and Rachels 2011. It is more controversial to ascribe implicit appeals to systematic normative frameworks, but I think that certain of the arguments in Singer 1977 and Norcross 2004, for example, make the most sense if we presuppose the consequentialist framework that the authors of these articles accept.

³ Many of the moves in the argument of this section and the next are familiar in the literature. They overlap most substantially with parts of Rachels 1997, DeGrazia 1996 and 2009, and McMahan 2008.

⁴ Harman 1977 uses a similar example for very different purposes.

conference with a kitten, gasoline, and matches, and announced my intention of testing philosophers' intuitive moral judgments with a real-life case. I would be shocked if my colleagues were not moved *en masse* to stop me.

Supposing that Cat is true, what is the best explanation of this fact? Some views suggest that one can explain the truth of Cat without suggesting that animal suffering is really what matters morally. Consider two such views. Immanuel Kant suggests that cruelty to animals is wrong because it is "...demeaning to ourselves" (1997, 27:710 [434]). Another view suggests that being cruel to animals is objectionable because it tends to make you callous, and thus apt to be cruel to humans.⁵

Perhaps cruelty to animals is typically objectionable in part for these reasons. However, these explanations cannot explain the full range of our relevant intuitive judgments involving Cat.⁶ Consider one example. Imagine a video game – *Cat Torturer!* – in which players score points by playing at torturing cats in a variety of increasingly nasty and vividly rendered ways. As I understand Kant, he would take playing *Cat Torturer!* for enjoyment to be demeaning to oneself. Further, consistently playing *Cat Torturer!* might well foster callousness in players towards the suffering of humans. However, there appears to be a striking moral difference between playing *Cat Torturer!* and actually pouring gasoline on a cat and lighting it on fire. The latter action appears worse exactly because it causes horrendous suffering to an actual cat. This example suggests that the views just mentioned ignore part of the best explanation for the truth of Cat. This is that animal suffering is intrinsically morally significant in a way that can explain why causing an animal to suffer is morally wrong. The rest of this section develops the case for this explanation.

My explanation appeals to suffering, which is a complex phenomenon. I take it that humans can suffer in ways that non-human animals cannot. However, I take it that it is plausible that humans and the animals I discuss are both capable of what I will call *visceral suffering*: being in the sort of pain caused by physical injury, and being averse to that state. I will assume the slightly controversial thesis that animals are capable of visceral suffering. (I return to this assumption in §4.)

⁵ The empirical case for this latter suggestion is unclear. For a brief discussion, see Herzog 2010, 31-7, and the papers cited there.

⁶ For a quite different, *Kantian* case against Kant's stated view, see Korsgaard 2004.

The claim that animal suffering is intrinsically morally significant fits smoothly with an attractive general account of the normative significance of suffering. Imagine that you are doing some amateur carpentry, and pause to consider the possibility of smashing one of your fingers with a hammer. Suffering the throbbing pain of an injured finger would typically be *intrinsically bad* for you. What this state would be like gives you *reasons* to hammer cautiously. And it helps to explain why it would be *wrong* for me to smash your fingers.

These explanations are not *exhaustive*, for at least two sorts of reasons. First, the wrongness of an instance of causing suffering can be outweighed or undercut by other features of that instance. For example, consider a dentist performing a painful but needed root canal on a consenting patient, or a coach pushing an athlete through a grueling training regimen. Second, my smashing your fingers may typically be morally objectionable for reasons beyond the suffering it causes. For example, it may express disrespect for your person, or some other vicious attitude or trait, and it may interfere objectionably with your agency. My claim, compatible with these points, is that causing visceral suffering is typically sufficient to explain the wrongness of an act.

To see this, compare two ways to express disrespect for someone, and interfere with his agency. On the one hand, you may inflict significant pain on him. On the other, you may insult him repeatedly in a loud and distracting way. Suppose that such insult is disrespectful and interferes with his agency, but does not cause him to suffer. If these were your two morally best choices (for example, because otherwise, Williams' malevolent militia captain executes both of you) you would surely typically be required to insult rather than inflict suffering. The most natural explanation of this fact is that the way it feels to suffer can suffice to explain the wrongness of inflicting significant suffering, independently of considerations of disrespect and interference (compare Rachels 2011, 883-4).

This plausible explanation of the wrongness of inflicting visceral suffering makes it very difficult to avoid the conclusion that it is typically wrong to inflict visceral suffering on animals. This is because the very feature that suffices to (nonexhaustively) explain the wrongness of causing visceral suffering in humans (the way such suffering feels) is by hypothesis present in cases of inflicting suffering on animals. This explanation also helps to explain our thinking about a range of cases involving animals. Thus, many of us think that

practices like dogfighting and cockfighting are wrong. Similarly, many people who take it to be permissible to kill and eat animals nonetheless take the conditions of contemporary factory farming to be morally problematic. The central feature of factory farms that people find morally horrifying is the suffering that such farms inflict upon animals.⁷

This section began with a plausible intuitive judgment – Cat – and argued that the best explanation of the truth of this judgment is that the way it feels to suffer can adequately (but not exhaustively) explain why it is wrong to cause suffering. The next section builds on this conclusion to complete my intuitive case for ethical veganism.

2. The intuitive-explanatory case for veganism

In this section, I argue for veganism in two stages. First, I argue for the wrongness of killing animals. Then, I argue that this entails that we should adopt a vegan lifestyle.

One might accept that it is wrong to cause animals to suffer, but deny that is wrong to kill animals. This combination of views underwrites the position of one sort of *ethical omnivore* (compare e.g. Pollan 2006, Ch. 17). This ethical omnivore argues that it is objectionable to consume many of the animal products that we do consume. However, this is because those products are produced by institutions that cause egregious animal suffering. Find a farm that allows a cow to graze, rather than restricting its movement in a tiny stall, and that slaughters its animals humanely rather than in a terrifying and painful manner, and it is morally acceptable to drink the milk and eat the meat that comes from that farm.

The resulting view can appear intuitively attractive. On the one hand, by insisting that animal suffering is intrinsically morally objectionable, the ethical omnivore can accept the case for Cat suggested in the previous section. It can also explain why the standard methods of raising veal are objectionable,⁸ why we should condemn slaughterhouse dismemberment of live animals, leg-hold traps, dogfighting, and *foie gras*. On the other hand, by denying that animal death is intrinsically morally objectionable, the view avoids

⁷ For a brief summary of some of the relevant facts, see Singer and Mason 2007, Ch.4.

⁸ Veal calves are typically raised in pens designed to be so small that they are effectively prevented from moving and thus from developing muscle tone, and they are also given a diet that purposefully malnourishes them. Together, these factors lead to their short lives being characterized by manifest distress, injury and sickness.

banning all animal products from our tables. Further, our judgments about the wrongness of killing animals are in general weaker and less clear than our judgments about causing animal suffering, so it can seem that there is a small intuitive price to pay for these attractive results.

Despite these attractions, I think that underlying premise of the ethical omnivore's view is indefensible. This is the claim that while animal suffering is morally objectionable, painless animal death is not. The difficulty with this claim can be brought out by considering a pair of cases, and developing a plausible partial account of the wrongness of killing that is the best explanation of our judgments about these cases.

First, suppose that one has an inspired idea for an art film. The film, however, would require performing a painful and unnecessary medical operation on a cow. If the suffering of animals is objectionable (as I have argued and our ethical omnivore grants), then there is a substantial moral objection to producing this film. Second, suppose that in order to save a different cow's life, one would need to perform an equally painful operation on the cow. If the cow would go on to have a long and flourishing cow life after the operation, this operation seems to be wholly unobjectionable and perhaps morally laudable.

This pair of cases shows something important: it can be (at least) permissible to cause suffering to an animal exactly because doing so is necessary to save its life. This suggests that the life of an animal is morally important: saving it can suffice to make an act that is typically impermissible (causing the animal to suffer) permissible. This point is clear, and very difficult to plausibly square with the claim that it is typically morally permissible to kill an animal. If killing an animal is unobjectionable, then why should saving its life give you sufficient reason to cause otherwise wrongful suffering (compare McMahan 2008, 67)?

Further, the best explanation of our judgments about this thought-experiment is provided by a plausible partial account of the wrongness of killing. Just as with the question of why it is wrong to cause humans to suffer, the correct story of why it is wrong to kill us will involve many elements. For example, killing a person typically interferes with their autonomy, understood here as their ability to choose and live the life that they value, to the best of their abilities. Killing is also typically inconsistent with adequate *respect* for that autonomy: it interferes radically with the victim's life in a way that they have not consented

to. If we assume that animals do not have autonomous plans, killing them painlessly is not objectionable in these ways.

There is another important reason why it is typically wrong to kill a person: killing typically deprives the victim of an objectively valuable future.⁹ That is, killing someone deprives them of the valuable experiences activities, projects, etc. that they would otherwise have had. The force of this explanation of the wrongness of killing can be brought out by considering cases of *life-extending killings* (Lippert-Rasmussen 2001). For example, suppose that there is a drug that, if taken, is known to damage one's heart such that one dies quickly, painlessly, and unavoidably, a year after ingestion. Ordinarily, giving you such a drug would simply be wrongful killing. However, suppose next that this drug is the only antidote to a poison that you have just accidentally ingested, which will otherwise kill you within the hour. Suppose finally that I administer the drug to you while you are unconscious from the poison, and sure enough, it extends your life, but causes you to die of heart failure a year later.

In this case the drug that I administered is the cause of your death in one clear sense. The coroner, in explaining why you died, would correctly point to the drug that I administered. Giving you the drug caused your death, and hence killed you. And yet, it simultaneously neutralized the poison and hence extended your life. This does not seem to be a case of wrongful killing. The crucial difference between this case and the ordinary case of life-shortening killing is that in the latter case, administering the drug deprives you of the presumably valuable future that you otherwise would have had, while in the former case, it increases the valuable future available to you. The best explanation for why life-extending killing in such a case seems distinctively unobjectionable is that an important part of what makes killing wrong is that it deprives the victim of a valuable future.

One might object that in this case my beneficent intentions do the crucial explanatory work. I agree that intentions may help to explain the wrongness of some actions, but this does not undermine the point made here. To see this, consider two cases where I give you the drug: an ordinary life-shortening poisoning, and the life-extending case just described. Suppose that in both cases I only administer the drug because I have

⁹ Readers will be reminded of Marquis 1989. Marquis suggests that deprivation of a valuable future is the 'primary' thing that makes killing a person wrong. Because the autonomy-based concerns sketched just above also strike me as crucial to a full account of the wrongfulness of killing, I am skeptical of this strong claim.

made a large bet that you will die exactly a year from today, and wish to collect. In the life-extending case, I would still be doing the right thing, just for the wrong reason.

In §1 I argued that suffering makes the lives of animals go badly. I take the complement of this claim to be similarly plausible: other features can make those lives objectively good for animals to have. If part of the explanation of why it is wrong to kill a person is that it deprives the victim of a valuable future, and animals can have valuable futures, it would be perplexing if depriving animals of such futures could not similarly explain the wrongness of killing them. Features that might make animal futures valuable are not hard to sketch. Animals seem capable of pleasures as well as pains, and pleasant lives are typically better. Similarly, healthy animals typically have better lives than unhealthy animals, pack animals lead better lives if they have companions, etc.¹⁰ Consider a range of things that one might do to an animal: isolating it, amputating a healthy limb; purposely raising it on a diet lacking essential nutrients, etc. These sorts of acts seem wrong. A natural explanation is that they are wrong in part precisely because they worsen the future available to the animal.¹¹

One might object that having a valuable future depends upon one's richly *valuing* one's future, and that very few animals value their own futures in the relevant sense. However, the value of an activity or state for a person does not appear to depend either on their eventually valuing it, or on their being psychologically capable of valuing it (compare Marquis 1989, 195ff for relevant discussion). Consider an unswerving misogynist, dismissive of the contribution that his relationship with his wife makes to his life. This sort of blindness may make his life worse, but it needn't erase the good constituted by the underlying relationship. If animals are incapable of valuing, this entails at most that they are in a situation analogous to that of such people.

Further, the explanation that part of why killing is wrong is that it deprives the victim of a valuable future fits beautifully with the cow operations thought-experiments. In

¹⁰ One should not conflate the most valuable life for an animal with the life that the animal tends to have in its 'natural' evolutionary environment. Some species exemplify a high attrition pattern, bearing many young each breeding cycle, few of whom ordinarily survive to adulthood. This does not entail that dying young is part of the good life for an individual member of such a species.

¹¹ DeGrazia 2009, 161-2 expresses sympathy for both hedonistic and functional characterizations of the positive value of animal lives. He is, however, cautious concerning the relative ethical significance of painlessly killing animals. I take the cow operations cases to ameliorate the need for such caution.

the second thought-experiment, it is the preservation of a valuable cow future that explains why it is permissible to cause otherwise suffering in a way that would otherwise be wrongful. If the explanation that I have offered is right, then the ethical omnivore's crucial premise – that while causing animal suffering is objectionable, causing animal death is not – must be rejected.

The explanation of the wrongness of killing that I have offered might suggest a different defense of the ethical omnivore. The defense points out that a pig raised for meat on a sufficiently humane farm may have a life that is on balance worth living. And were it not raised for meat, the pig would not exist at all. So, the ethical omnivore might suggest, raising the pig for meat is part of a plan that is good for the pig, and bad for no one. How could it be wrong?¹² This defense raises serious theoretical controversies, usually discussed under the rubric of the 'repugnant conclusion'.¹³ However, we can set these controversies aside, by focusing on an analogous case which does not involve the creation of new life.

Suppose that a dog will be put down by the animal shelter unless Al adopts it. And suppose that Al only wants a dog in order to perform a painful and unnecessary operation on it, for his art film. Suppose finally that Al will be careful to give the dog a long and good life, such that he is reasonably sure that the dog's life in his care will have been worth living, despite the pain and degraded capability produced by the operation. It seems to me that Al would act wrongly in performing the operation, despite the fact that this is part of an overall pattern of his action that has benefitted the dog. In general, no amount of good treatment of an animal can license one to substantially harm that animal simply for one's own ends.¹⁴ But this is precisely what the 'ethical' pig farm by hypothesis does: it gives the pig a good life only to radically shorten that life by killing it, in order to serve the interests of the farmers. As I have just argued that killing an animal young harms the animal by depriving it of a valuable future, the ethical pig farm is morally objectionable for the same reason that Al's operation would be.

¹² This objection is strongest when posed on behalf of the *ethical* omnivore. For one might reasonably deny (as Rachels 2011, 884 does) that factory farmed animals have lives worth living.

¹³ The classic text is Parfit 1984, §131. For a review of the extensive literature, see Arrhenius et al 2010.

¹⁴ Compare DeGrazia 2009, 162. While DeGrazia talks of 'unnecessary' harm, I think that we substantively agree: I do not think that the fact that Al *would not* adopt the dog without performing the operation makes doing so necessary in the sense that DeGrazia has in mind.

If what I have argued to this point is correct, institutions that cause widespread animal suffering or death are thereby engaging in systematic wrongdoing. I claim that this conclusion supports adopting a vegan lifestyle. This is because there are constraints on how we may permissibly interact with morally objectionable institutions. This idea may be motivated by example: if one knows that a certain bar of chocolate is produced using child slave labor (as much of it is),¹⁵ it is very natural to take the purchase of that chocolate to be objectionable. The explanation of why will be both complex and controversial, and I will not do it justice here.¹⁶ However, an initial gloss would begin with the fact that in knowingly purchasing and enjoying the chocolate, you would be:

- (a) *seeking to benefit from* the wrongful acts of an institution (cocoa plantation child slavery), by enjoying the fruits of its wrongful activity, and
- (b) *cooperating with the very plan* that rationalizes the wrongful acts (cocoa slavers hope to gain from producing cocoa by using child slaves exactly because they rely on the willingness of consumers to purchase products made with this cocoa, whether out of indifference or ignorance).

If this is the right preliminary explanation of the wrongness of knowingly purchasing slave chocolate, it will also apply to purchasing animal products produced by institutions that treat animals in seriously objectionable ways.

One might think that the argument thus far is compatible with being an *ethical vegetarian*, who consumes animal products, but not animals, rather than a vegan. The leading thought suggesting such compatibility can be illustrated by an example. Producing milk is a normal and constant part of the life of cows. Indeed, cows have been bred to be

¹⁵ Most of the world's cocoa is produced in West Africa, on farms that employ large numbers of child laborers working in horrendous conditions, a significant proportion of whom are victims of human trafficking. For one brief journalistic treatment, see Orr 2006.

¹⁶ Some philosophers content themselves with the thought that the conclusion is deeply plausible, without seeking a deep explanation (compare Rachels 1997). Many alternative models of explanation are possible here. On an 'expressive' view, the well-informed purchaser of the chocolate could be claimed to objectionably express acceptance of the oppression of the slave (compare Anderson 1993, Ch. 2). Alternatively, one might attempt to adapt to this case Scanlon's proposal (2008, Ch. 4) that we should 'downgrade' certain aspects of our relationships with wrongdoers. Finally, one might approach the problem in a consequentialist way. For example, Norcross 2004, 232-3 argues that, although the food system is not perfectly sensitive to consumer demand, if some large number of people ceased to buy animal products, there would be a marked decrease in animal killing. In such a pattern, each individual vegan has a tiny chance of being the individual whose choice happens to be causally efficacious. But if she is causally efficacious, she will save a huge number of animals at a stroke. A tiny chance of making such a large difference is morally important, according to Norcross. For an influential and structurally identical treatment of voting, see Parfit 1984, 735; for an important argument against using this consequentialist reasoning to defend vegetarianism, see Budolfson *ms*. Finally, see Kutz 2000 for a detailed discussion of the ethical significance of relationships to objectionable institutions.

such that failing to milk them would be intensely cruel. It may seem that there could be no objection to drinking this milk, if the producing cows are treated well. After all, one does not harm the cows by taking their milk.

While I take this point to be correct in the abstract, its application ignores the economic realities of (even ‘humane’) contemporary dairy farming. Robust milk production requires roughly annual calving. The majority of the calves (and almost all of the male calves) are typically raised to be killed early, and failing to raise them in this way would constitute an enormous drain on a farm’s resources. The economic logic of dairy farming thus requires the very early killing of most of the animals involved. But if killing animals is wrong for the sorts of reasons that I have sketched, then almost all existing ‘humane’ dairy farms systematically engage in the wrongful treatment of animals. (Note that this explanation rests on facts about actual farming practices. It thus allows that there are conceivable farming practices whose animal products it would be unobjectionable to consume.)

The argument of this section and the last began with the intuitive claim that it is wrong to set fire to a cat, and argued that once we accept that claim, it is very hard to resist an intuitive case for ethical veganism. This case entails that not only the typical North American diet, but also ‘ethical omnivorism’ and ‘ethical vegetarianism’ are typically morally objectionable. These are strikingly revisionary conclusions. In the next section, I show that this fact suggests an important methodological challenge to this intuitive case.

3. The methodological challenge: one person’s modus ponens...

In this section, I first argue that the case developed in the previous two sections can be plausibly re-interpreted as defending a conditional claim: that if Cat is true, we should accept ethical veganism. If true, this conditional claim forces us to revise our intuitive views about animal ethics. However, it is possible to challenge the conclusion that I draw from this argument, by questioning whether we must revise our intuitions in a pro-animal direction. (By *pro-animal*, I mean a view or judgment that requires us to treat animals well in some respect. I will also use *anti-animal* to describe views that do not require us to treat animals well).

Call an ethical account about a topic that vindicates our clear reflective intuitive judgments about the topic a *conservative* theory. Call the view that some conservative theory is defensible in animal ethics the *conservative view*. Conversely, call an account that requires that we abandon some intuitive judgment about a topic a *revisionary* theory, and the view that such a theory is required in animal ethics the *revisionary view*. Where they are available, conservative theories are *prima facie* highly attractive. After all, conservative theories are theories that are consistent with all of our clear reflective judgments about the subject matter. Revisionary theories, by contrast, require us to abandon something that we believe. They thus undertake a serious explanatory burden: to explain why we should dismiss the abandoned judgments.

The first thing to notice about the argument for ethical veganism of §§1-2 is that its force rests largely on the plausibility of intuitive judgments like Cat. If we rejected Cat and related intuitive judgments, the whole case that follows would be cast into doubt. To see this, consider an important competitor to my explanation of the wrongness of causing humans to suffer offered in §1. On this view, the wrongness of causing human suffering is explained by what it is like to suffer, together with the distinctive ‘moral status’ of the suffering human, where this status is claimed to be an essential feature of every human, whatever her individual capacities (compare e.g. Cohen 1986, 866). One important (I think near-decisive) problem for this sort of view is that it makes it very difficult to offer the most natural explanation of the wrongness of torturing animals for fun, dogfighting tournaments, dismembering live animals in slaughterhouses, etc. However, for someone initially unmoved by Cat, and these related cases, this would not appear to be a cost of the ‘moral status’ explanation.

This has an important consequence. It suggests that my argument of §§1-2 can be partially but accurately represented as follows:

Cat	It is wrong to pour gasoline on a cat and light it on fire
Conditional	If it is wrong to pour gasoline on a cat and light it on fire, then (inter alia) it is wrong to drink a glass of milk
Not-Milk	It is wrong to drink a glass of milk

This representation follows from the suggestion just made that my argument can be separated into an intuitive appeal to Cat, and an argument for Conditional. But this point

suggests an important challenge: why cannot one simply offer the ‘modus tollens’ version of this argument?

Milk	It is not wrong to drink a glass of milk
Conditional	If it is wrong to pour gasoline on a cat and light it on fire, then (inter alia) it is wrong to drink a glass of milk
Not-Cat	It is not wrong to pour gasoline on a cat and light it on fire

Like Cat, Milk has been chosen to serve as a vivid proxy for a cluster of related and (for most of us) highly plausible judgments.¹⁷ Cat stands in for judgments espousing the wrongness of engaging in various acts of cruelty to animals. Milk stands in for judgments espousing the permissibility of our ordinary everyday use of animal products.

The availability of the modus tollens version of the argument suggests that (even if otherwise sound) what my argument from §1-2 most clearly shows is that the conservative view of animal ethics is indefensible: we must give up either Milk (and related judgments) or Cat (and related judgments).

It is worth noting that even if the case developed in §§1-2 is best understood as an argument against conservatism, it answers an important sort of worry that one might have about ethical veganism. This worry is that philosophical reasoning may not be epistemically powerful enough to successfully challenge something as deeply embedded in our practices as using animal products. Using animal products has been taken for granted as permissible by virtually the entirety of humanity throughout history. Could philosophical reasoning really entitle us to the conclusion that such an activity is morally wrong? One can partially answer this worry by pointing to the manifest soundness of various revisionary arguments offered against deeply engrained racism, sexism or homophobia. However, the case against conservatism suggests another retort: it shows that *any* defensible set of ethical beliefs about animals will have to abandon some highly intuitive theses, on pain of inconsistency. Ethical veganism is revisionary, to be sure, but so is the view that it would have been acceptable for me to test your intuitions by actually setting fire to a cat. If Cat and Milk cannot both be maintained, ethical veganism does not appear to be vulnerable simply in virtue of its revisionary character. Rather, from this point on, we need to assess the case

¹⁷ Note that I say *most of us*. Some people presumably think that there is nothing wrong with setting fire to a cat, and a good deal more think that there *is* something wrong with drinking milk. However, I take both groups to be well outside of the intuitive mainstream.

for retaining Cat *rather than* Milk, (or vice-versa). If she can force us to this choice point, the ethical vegan has made crucial dialectical progress.

Call the thesis that we must give up either Milk (and related judgments) or Cat (and related judgments) *The Tension*. (In what follows I will sometimes describe Cat and Milk as being ‘inconsistent’ or ‘in conflict’. These should always be read as references to *The Tension*, and not as claims about meaning or logical form). This problem is not unique to my argument for veganism. Rather, I suspect that it will apply to any revisionary ethical argument that focuses on developing plausible ethical explanations grounded by intuitively powerful judgments (what I earlier called the ‘intuitive-explanatory’ method). The fact that some of the strongest pro-animal arguments have this form raises a pressing methodological question: how should we adjudicate this sort of conflict in our intuitive judgments? I now turn to that question.

4. The methodological implications of revisionism

I have just argued that the failure of the conservative position leaves us with an important and neglected question: should we resolve *The Tension* by abandoning Cat (and related judgments), or Milk (and related judgments)? Backing up a little, we can ask a methodological question: what is a philosophically principled way of resolving *The Tension*? This section is dedicated to addressing the second, methodological question. I will briefly consider and reject three salient strategies for resolving *The Tension*. The first is to appeal to the greater intuitive plausibility of one over the other resolution. The second strategy is to appeal to systematic normative ethical theorizing as a route to resolving *The Tension*. Finally, the third strategy seeks to resolve *The Tension* by defeating a non-moral presupposition of our belief in Cat. I suggest reasons for pessimism about all of these, arguing that the only philosophically satisfying resolution to *The Tension* requires appealing to a debunking argument that undercuts the intuitive plausibility of one of the inconsistent clusters of claims.

The first strategy is to attempt to ascertain which of the inconsistent claims can be abandoned with the least reflective implausibility. For example, perhaps with the case for *Conditional* in place, the ‘modus ponens’ case against Milk just seems more plausible than

the ‘modus tollens’ case against Cat (or vice versa). One initial worry about this variant of the strategy is the difficulty of developing a persuasive argument against someone inclined on the same basis to resolve The Tension in the other direction.

The intuitive strategy can be developed in two more interesting ways. First, this strategy can be developed in the spirit of G. E. Moore’s notorious arguments against the skeptic and idealist. To begin, notice that Conditional is supported by the conjunction of a whole series of intuitive and explanatory claims. It collapses if any one of them is rejected. One might then ask, in a Moorean spirit: am I *more certain* in the conjunction of Cat and Milk, or in the complicated conjunction that underwrites Conditional? One difficulty with this strategy is that it threatens to prove too much. After all, this form of argument, if legitimate, would seem to threaten almost any multi-premise intuitive-explanatory philosophical argument for a revisionary conclusion. A diagnosis of the difficulty is that Moore’s strategy should not be used to support just any plausible claim. Rather, Moore typically appealed to the *most plausible* deliverances of common sense, like “I have a hand”, or “things move”. It is implausible that either Cat or Milk deserve a like status as “Moorean facts”.¹⁸

A more promising variant of the intuitive strategy would seek to show that one of the conflicting sets of intuitive judgments is linked by an explanatory structure to a much wider set of our judgments than the other. If one could show this, then one would have a principled way of defending one resolution of The Tension over the other. However, I am not sure how to make such a strategy compelling, for reasons best illustrated by example.

In one of the few discussions of animal ethics that notice roughly the ‘modus tollens’ problem, Alastair Norcross deploys this strategy.¹⁹ Norcross argues that we should retain the pro-animal judgments because this permits us to offer the most plausible account of our obligations to ‘marginal’ humans (humans who lack robust rational capacities). However, this part of Norcross’ argument faces two structural difficulties. The first problem is that ‘marginal’ humans are a notoriously hard ethical case: it is extremely

¹⁸ See my 2009 for a detailed discussion of the use of Moorean arguments in ethics, that supports the points made here.

¹⁹ 2004, §§3-6. Norcross’s version of the competing arguments also stacks the deck, taking the premise “it is not wrong to support factory farming” as his opponent’s starting point. But *that* is wildly less intuitive than the claim about milk that I consider.

difficult to develop an explanatorily satisfying theory in this area. It is thus unclear how costly it is to fail to provide a satisfying theory of these cases. Second, Norcross' own theory also makes highly counterintuitive claims about these cases. For example, it entails that I would do something roughly as seriously wrong by setting fire to a cat as I would by setting fire to my infant son.

I take Norcross to have chosen the most promising case for his strategy. 'Marginal' humans are arguably distinctively relevant to animal ethics, because some marginal humans have capacities that are at least roughly analogous to those of relevant non-human animals. However, because the most promising analogous cases are independently so vexed, it appears unlikely that this strategy can bear fruit.

Note next that, as I have set it out, our inconsistent set is explanatorily lopsided: in developing my case for *The Tension*, I basically began with *Cat*, and then spelled out an explanatory case for ethical veganism. By contrast, it may not be clear what an explanatorily satisfying argument against *Cat* would look like. However, this is an artifact of my presentation. There are a variety of systematic normative theories that imply that we are not required to treat animals well. These constitute candidate explanations of the falsity of *Cat*.

This might suggest that we should resolve *The Tension* by looking to systematic normative theory. This approach is well-worn in the literature. For example, Peter Singer's pro-animal view (1977) is arguably most compelling when embedded within his utilitarian framework, while William Baxter's case that animals only matter if and when we care about them (1974) is grounded in a simple contractarian ethic. Defending such global theories is admittedly hard, but the broader explanatory power promised by such theories might appear to be a way to make progress when forced to choose between revisionary claims in a case like ours.

While I am a cheerleader for systematic normative theorizing, I worry that here the apparent aid offered by such theorizing is probably illusory. To see why, consider classical utilitarianism: this consists in a theory about what value consists in (a positive balance of pleasure vs. pain), and a structural theory about how we should respond to value (by performing acts that maximize it, wherever it might be located). Such a theory comes close to forcing a pro-animal conclusion upon us. However, many of the heirs of classical

utilitarianism have tended to find only its consequentialist maximizing structure distinctively compelling. From J. S. Mill onward, many have tended to abandon simple hedonism about value. Instead, they have adopted pluralistic theories of value, largely driven by concern to capture our intuitive judgments about value. However, it is simple hedonism about value, and not consequentialism, that is essential to the classical utilitarian case in defense of animals. To see this, note that a consequentialist might argue that only the pleasures and pains of beings with a *moral status* lacked by most non-human animals should enter into the maximizing calculus. The natural way to determine whether this theory is more or less plausible than classical utilitarianism is by appealing to the very intuitive judgments about animals that The Tension has cast into doubt. If this is so, we should not expect appeal to utilitarian theoretical structure to help us to independently resolve The Tension.

The same sensitivity to local intuitive judgments cuts against the force of seemingly animal-unfriendly systematic theories. To see this, consider the work of T. M. Scanlon, arguably the leading contemporary contractualist. Scanlon is careful to argue that his theory can be adapted (either via restricting its scope or via trusteeship hypotheses) to protect animals (1998, 177-184). Again, the structure of the theory is left intact, while its distinctive force to adjudicate our debate about animals has been basically eliminated.

These examples suggest two points. First, pro- or anti-animal modules can typically be grafted onto the central structure of most important normative theories. Second, the main reasons for accepting or rejecting those grafts will arise from their *local* intuitive-explanatory plausibility: how well they vindicate our careful thinking about our ethical relationships to non-human animals. If this is true, then it is not clear that appeal to systematic normative theory is a promising way to adjudicate The Tension.

In discussing these examples, I have presupposed the arguably dominant mode of normative theorizing: one that gives a distinctively privileged and robust methodological role to our intuitive moral judgments. There are important methodological alternatives in normative ethics, typically grounded in broadly metaethical commitments. Consider two examples. First, Christine Korsgaard advocates a methodology in ethics that involves a kind of 'practical conceptual analysis' that is allegedly called for by the practical problem that we confront (2003, 115-6). Second, Richard Boyd (1997) develops a metaethic that diminishes

the usual epistemic role of intuitive judgments, and suggests that some seemingly intractable moral conflict may be explained by the presence of indeterminacy in the moral facts. Discussing the methodological significance of such theories carefully would exceed the scope of this paper, so I set them aside here (or discussion, see my 2012).

A third possible strategy is to seek to resolve The Tension by defeating a non-moral presupposition of one of the inconsistent theses. A version of this strategy that attacks the presupposition that animals can suffer is the heart of perhaps the most familiar anti-animal philosophical strategy. Historically, the claim that animals can suffer has been most famously challenged by René Descartes. On Descartes' view, animals are just complicated machines with no inner lives: just as there is nothing that it would be like to be a pulley or a lever, there is nothing that it would be like to be a cow.²⁰ More recently, the assumption of animal suffering has been most carefully challenged by Peter Carruthers. Carruthers argued that while animals can perhaps have pain, they do not have *conscious experience* of pain, which he argues is the morally relevant property (1992, Ch. 8).²¹ These claims promise to ground a revisionary anti-animal view, because they would seem to entail that setting fire to a cat is no more objectionable than setting fire to a bicycle.

The presupposition-defeating strategy is attractive because it avoids the potential dialectical stalemate threatened under the previous two strategies. If Carruthers' claims about animal capacities could be defended,²² The Tension could seemingly be resolved 'from the outside'. This is because the plausibility of Cat would not survive conviction that there is nothing that it is like to be a cat.²³

This very feature of the presupposition-defeating strategy should leave us philosophically dissatisfied. The problem is that convincing us that animals cannot suffer resolves The Tension, but leaves the deeper intuitive inconsistency intact. To see this, note

²⁰ For example, see his letter to Mersenne, 1991 [1640], 148. For an interpretation that challenges elements of this standard reading of Descartes on animal experience, see Thomas 2006.

²¹ Carruthers has since suggested (in his 2004 and elsewhere) that animals may count as suffering in virtue of finding their pains awful, even if they lack phenomenal consciousness.

²² I take this task to be nearly hopeless, but that is an argument for another day.

²³ This oversimplifies. For example, if the newer Carruthers view mentioned in n. 21 above were correct, animal suffering can matter even if it is not conscious. This shows that in order to resolve our inconsistency, empirical claims about animal capacities will still need to be wedded to controversial normative claims; in this case, claims about what exactly normatively significant suffering requires.

that prior to being convinced by Carruthers (or whomever) we would presumably accept not just Cat and Milk, but the following conditional claims:

- Cat*** If animals were capable of suffering, then it would be wrong to pour gasoline on a cat and light it on fire
- Milk*** If animals were capable of suffering, it would not be wrong to drink a glass of milk

My argument for Conditional in §1-2 can be adapted to show these to be inconsistent just as Cat and Milk are. Becoming convinced that animals cannot suffer may defeat our belief in Cat, and hence resolve the claimed ethical inconsistency between Cat and Milk. However, it leaves the deeper inconsistency – between Cat* and Milk* – intact.

Philosophical ethics does not merely aim to tell us what to do in practical cases. It aims to *explain* why we should do those things. Because of this explanatory ambition, an animal ethics that leaves this deeper inconsistency untouched is philosophically unsatisfying. Thus, it seems to me that the presupposition-defeating strategy (like the Moorean strategy mentioned above) is crucially philosophically incomplete. The intuitive and normative theoretic strategies discussed earlier in this section do not have this problem. But they have a related shortcoming. They take our incoherent judgments at face value, and seek to move from those judgments to a more coherent set. But they do nothing to explain *why* those judgments are incoherent in the first place. This point partially explains why on both of these views, there is basically nothing one can say to someone who has thought hard about the cases, and come to the opposite conclusion than you have. Both of you have successfully moved to a more coherent set of beliefs.

We can improve on this state of affairs. The discussion of §§1-3 suggests that we cannot have a conservative view that avoids The Tension. If this is so, we should aim to develop an account that explains the existence of The Tension. What would such an explanation look like? Given the nature of The Tension, I think that it can only take the form of a *debunking argument*, which explains why we are illicitly tempted by either Milk or Cat. Let me very briefly say what I mean by a debunking argument. I will follow Guy Kahane (2011) in suggesting that debunking arguments provide alleged undermining defeaters for the beliefs that they target. That is, rather than providing positive evidence that a belief is false, they aim to discredit the evidence that we take ourselves to have for the truth of that belief. If the debunking argument for abandoning one of our two

inconsistent judgments is much stronger than that for the other, this gives us a principled way of resolving The Tension. In the next section, I argue that this is the case.

5. Debunking intuitive judgments about animals

In this section, I set out the most plausible debunking argument available to cast doubt on our belief in Cat and Milk respectively. The debunking argument for belief in Cat that I consider appeals to the idea that this belief arises from illegitimate anthropomorphization. The argument for debunking Milk argues that this claim arises due to a kind of rationalization. I will argue that the most plausible debunking argument for Cat is substantially weaker than that for Milk, and that this gives us grounds to prefer the pro-animal resolution of The Tension.

Debunking arguments have a controversial place in ethical theorizing. On the one hand, some of the most important figure in the history of ethics and political philosophy – like Karl Marx and Friedrich Nietzsche – can be understood as partly attempting to debunk some of the prevailing ethical views of their day. Further, whatever you think of these figures, it is hard to avoid the conclusion that across large stretches of human history, many prevailing beliefs about the ethical status of slavery, political organization, race, class, gender, and sexual orientation are reasonably described as the products of ideology, in need of debunking. On the other hand, debunking arguments play a limited role in contemporary ethics, and this is perhaps in part because they can appear to be problematically blunt philosophical implements, too easy to deploy indiscriminately against anything one doesn't like. (As in the sophomoric parody: "Dude, you can't trust your moral judgments; you just believe them because The Man wants you to.")

There are thus very hard questions about where exactly debunking arguments have force, and how much force they have. However, if my argument in the paper thus far is granted, ours is a best-case scenario for deploying such arguments. This is for two reasons. First, I have argued that there is an inconsistency between two groups of our intuitive judgments (represented by Cat and Milk). Second, I have argued that an explanatorily fulfilling resolution of that inconsistency must include an explanation of why one of these groups of intuitive judgments is less credible than it initially appears. This is exactly the

context in which debunking arguments should be most powerful, because it is a situation where only such an argument can provide a view in animal ethics that helps us to fully understand *why* we should resolve the tension in our beliefs in one way or another.

I cannot consider every possible candidate debunking argument here. Instead, I will focus on what I take to be the single most plausible argument available to each side. First consider what I take to be the most plausible debunking argument on behalf of the anti-animal view. This is the thesis that belief in Cat, and related claims, are the result of a tendency to anthropomorphize companion animals. The debunker points out that if we exchanged Cat for a similar claim involving rats or mice, for example, the intuitive revulsion will be much lessened. Part of the typical culture of our relations to companion animals (and perhaps part of what tends to make those relationships so emotionally rewarding for the humans involved) is a tendency to treat the animals as if they had human-like psychologies. However, while it may be good for our emotional lives to treat animals this way, belief that they have such a psychology does not survive reflection. For example, the temptation to blame Fluffles for destroying your houseplants ‘in a fit of pique’ should recede when one asks oneself if it really makes sense to take the Strawsonian ‘participant attitude’ towards her behavior.²⁴ Given Fluffles’ nature, it may make sense to protect your remaining plants, or to attempt to train her, but it does not make sense to treat her as an agent apt for blame or punishment.

If our intuitive acceptance of Cat and similar pro-animal intuitions is explained by mechanisms that serve this anthropomorphizing tendency, the justificatory force of those intuitions is undercut. Rather than being responses to the moral significance of the relevant animals, these intuitions are being generated by emotional processes insensitive to such moral significance. This hypothesis predicts that our intuitions are less vulnerable to this sort of distortion when we consider animals (or especially types of animals) to which we have not formed emotional ties. In these cases, our intuitions about the significance of animal suffering are weaker. This in turn suggests that The Tension may be partly explained as a result of such anthropomorphization, and that the intuitive force of Cat should thus be discounted.

²⁴ Compare Strawson 1962. In conversation, some people have remained committed to their anthropomorphization, claiming, for example, that their dog knew *exactly* what she was doing, when she chewed up dad’s shoes to punish him for being late, and that she was blameworthy for doing so.

How effective is this debunking argument? There is a grain of truth in it, but only a grain. Our intuitions about causing suffering to companion animals are certainly typically clearer and more vivid than those concerning other animals. However, there is a plausible alternative explanation of the asymmetry that inverts the debunking argument. Consider the familiar idea that it is much easier to press a button, detonating a bomb that will kill far-off persons, than it is to bring oneself to kill a person with one's bare hands. There is a clear explanation for this: in the bare-handed case, the suffering of the victim – and its moral significance – is much more psychologically salient than in the button-pressing case. It would be wrongheaded to presume that our intuitions are much more reliable in the emotionally 'cool' button-pressing case than in the bare-handed case. This suggests a plausible competing explanation for the debunker's data. Our intuitions about the moral significance of the suffering of companion animals are typically stronger than our intuitions about animals whose lives are less familiar to us. However, this may simply be because in the former case, the morally significant suffering is more psychologically vivid to us, in light of our actual and imaginative familiarity with the animals in question.²⁵

The most promising debunking argument against our belief in Milk and related judgments argues that these are a product of what I will call *status quo rationalization*.²⁶ A belief is subject to status quo rationalization just in case that belief is required to vindicate the goals and behaviors of the believer, and of others that the believer identifies as members of her moral community. Consider as an example a member of a slave-owning family in the antebellum South, for whom owning and using slaves is a deeply embedded part of everyday life, and the life of those he is closest to. It is easy to predict that other things being equal, such a person will tend not to believe that slavery is a moral monstrosity. This is because taking oneself and those one identifies most closely with to be doing something seriously morally wrong makes for a particularly uncomfortable form of

²⁵ This explanation can grant *some* debunking of the processes that guide our intuitions: perhaps cuteness and ugliness play some role. For example, most of us are probably less apt to be sensitive to the suffering of star moles or pangolins than we are to that of baby seals or penguins. However, *these* sorts of differences can be adequately managed by the intuitive-explanatory method, exactly because they are unlikely to survive extended reflection.

²⁶ Christian Coons suggests another important debunking strategy to me in conversation. On my view, there is nothing intrinsically wrong with drinking milk (e.g. I suggested in §3 that there are possible morally unobjectionable dairies). Drinking milk is only wrong when and because it connects to the maltreatment of animals in the right way. The relevant version of Milk is thus that it is not wrong *to drink milk produced in ways that maltreat animals*. And there our intuitions may not be so clear.

cognitive dissonance: very few people can comfortably self-identify as morally bad. Sometimes this sort of cognitive dissonance is a catalyst for moral improvement. However, the more deeply embedded a behavior is in one's life – the more convenient or beneficial or pleasant; the more unquestioned by one's peers, etc. – the more likely that cognitive dissonance will be resolved by one's values changing to rationalize one's behavior, and that of those one identifies with.²⁷

Applied to the tension between Cat and Milk, the debunking argument goes like this: drinking milk and consuming other animal products is deeply embedded in the lives of almost all of our communities. Because of this, status quo rationalization is a highly salient explanation for why we take such consumption to be permissible (and hence, why we are inclined to accept Milk). On the other hand, singling out companion animals for torture is not a part of our culture in the same way. Because rationalization is apt to occur where it is needed to block obvious moral tensions, and because the moral tension between accepting Cat and our ordinary lives is not at all obvious, we would not predict such status quo rationalization to affect beliefs about setting fire to cats (except perhaps in philosophers who notice the tension).²⁸ This suggests that status quo rationalization is well-placed to explain The Tension: Milk but not Cat can be explained away as a product of status quo rationalization.

The availability of status quo rationalization as a potential explanation for an intuitive ethical judgment does not by itself debunk that judgment. Consider an example: the prohibition on wanton killing of humans is beneficial for me, and abandoning it would threaten some very powerful and central assumptions that my peer group holds about the moral asymmetry between themselves and serial killers, for example. The crucial contrast between this case and Milk is that in the wanton killing case, our moral beliefs are *overdetermined*. It may be true that status quo rationalization plays some role in bolstering our belief in the wrongness of wanton killing, but those beliefs also fit well with almost all

²⁷ Related rationalizing mechanisms presumably also underwrite the distorting tendencies of *partiality* in our moral thinking. For example, we tend to think that even someone who would attempt to be scrupulously fair should probably not sit in judgment on his own case. Compare Rawls 1951, 182 and Sinnott-Armstrong 2006, 195-7.

²⁸ Interestingly, Herzog 2010, 35-7 cites studies that suggest that childhood cruelty to animals is extremely common, admitted by a third to a half of all adults. However, Herzog suggests that such behavior may be understood by the actors themselves as transgressive, and this may explain why there is little pressure to rationalize it as permissible.

of our core moral reasoning. Milk is different in this respect. The argument of §§1-4 suggests that (a) Milk is part of a relatively explanatorily isolated set of judgments, and (b) Milk is (unlike the wanton killing judgment) in conflict with other plausible moral judgments. I suggested above that this is exactly the situation in which debunking arguments have the most dramatic force.

I have argued that the most plausible debunking explanation of the other inconsistent judgment – Cat – is weak. I have now suggested that the intuitions represented by Milk can be plausibly explained by appeal to the mechanism of status-quo rationalization. This suggests good grounds for taking status quo rationalization to be an undercutting defeater for the evidential status of Milk and related judgments. If this is right, the argument of this section has put us in a position to resolve The Tension identified in §3: debunking arguments give us some explanatorily satisfying grounds for preferring the case for veganism over anti-animal revisionary arguments.

Conclusions

In this paper, I have developed a case for a strongly revisionary view in animal ethics: ethical veganism. I began by setting out what I called an intuitive-explanatory case for this view. I suggested that this argument can be best understood as supporting a conditional argument, that connects accepting one plausible claim (Cat – that it is wrong to pour gasoline on a cat and light it on fire), to rejecting another (Milk – that it is not wrong to drink a glass of milk). The problem is that one can run a reverse argument retaining the same conditional thesis, but this time holding Milk fixed and concluding that Cat must be false. I then rejected three strategies for addressing the question of which direction of the argument should be preferred: attempting to adjudicate the intuitive force of the competing premises, appeal to systematic normative theories, and appeal to evidence against a crucial non-moral presupposition of one of the arguments. I argued that the only philosophically satisfying resolution of the conflict would appeal to an argument that debunks our belief in one of the crucial conflicting claims. Finally, I argued that one such argument – the claim that Milk is dubitable because a likely result of status quo

rationalization – is more powerful than the most salient alternative. I conclude that this gives us some reason to favor a strongly pro-animal view in animal ethics.

I want to close by emphasizing the modesty of my conclusion. First, I claim that considering relevant debunking arguments should lead us to *favor* a pro-animal resolution to this question. This is very different from saying that we have sufficient evidence to be justified in believing that it is wrong to drink milk, for example. If I am right, the most defensible substantive view about animal ethics supports ethical veganism *to a greater degree* than it does the status quo, or a variety of more modest pro-animal views. However, that does not entail that we have sufficient evidence to justify believing ethical veganism to be the ethical truth: I am inclined to think that given the difficulty of the case we may not. This does not, however, undercut the practical significance of this argument for ethical veganism. If the wrongness of an action is better supported by the evidence than its permissibility, this surely should lead reasonable deliberators to refrain from performing it, other things being equal.²⁹

Second, this argument does not address one of the main philosophical cases for the anti-animal position: if animals cannot suffer, then nothing that I have said here suggests that there is an objection to the anti-animal revisionist. As I argued in §4, the question of whether the empirical presuppositions of our intuitive ethical judgments are correct is independent of our best understanding about how to make those judgments coherent.

Nonetheless, I think the argument offered in this paper is important both substantively and methodologically. Substantively, the conclusion that veganism is better supported by the evidence than the alternatives *given* the assumption that animals can suffer is independently striking. It also highlights the ethical importance of empirical investigation of animal capacity to suffer. This is not because of the plausibility of the no-suffering view, which is minimal. Rather, it is because such investigation will inform how far the scope of the ethical vegan argument extends. Between oysters and cows is a continuum of animals with increasingly sophisticated cognitive capacities. The question of exactly which of those animals can suffer is an extremely difficult question that my argument suggests is strikingly ethically relevant.

²⁹ Here I make a modest point; in doing so, I seek to finesse the real complications that arise from the need to make decisions under moral uncertainty. See Ross 2006 and Sepielli 2009 for two discussions of this problem in its more general form.

Methodologically, I take my argument to provide an important model for defending revisionary views in applied ethics. Arguments for revisionary ethical views face strong dialectical burdens. Animal ethics is a case in point: in virtue of the counterintuitive conclusions of radical pro-animal arguments, it can always seem reasonable to reject them by rejecting one or another premise of such arguments. The worry is that the revisionist is by definition asking us to reject something highly plausible, so why not one of their premises! The argument of §§4-5 seeks to answer this worry. At least in cases with conditions similar to the animal ethics case (inter alia, a striking tension and detachability from structural claims in normative ethics), debunking arguments may have a more central role to play in applied ethics than is typically noticed.³⁰

Works Cited

- Anderson, Elizabeth, *Value in Ethics and Economics* (Cambridge, MA: Harvard UP, 1993).
- Arrhenius, Gustaf, Ryberg, Jesper, and Tannsjo, Torbjorn, 'The Repugnant Conclusion', *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*, ed. E. N. Zalta. <http://plato.stanford.edu/archives/fall2010/entries/repugnant-conclusion/>
- Baxter, William F, *People or Penguins* (New York: Columbia UP, 1974.)
- Boyd, Richard, 'How to be a Moral Realist', in S. Darwall, A. Gibbard, and P. Railton (eds.), *Moral Discourse and Practice* (New York: Oxford UP, 1997), pp. 105-136.
- Budolfson, Mark, 'The Ethics of the Marketplace and Some Surprisingly Deep Questions for Normative Theory,' *Ms.* Version 3.11. Available from the author.
- Carruthers, Peter, 'Suffering Without Subjectivity,' *Philosophical Studies* 121 (2004), pp. 99-125.
- , *The Animals Issue* (Cambridge, UK: Cambridge UP, 1992).
- Cohen, Carl, 'The Case for Using Animals in Biomedical Research,' *New England Journal of Medicine* 314 (1986), pp. 865-869.
- Cox, Christopher, 'Consider the Oyster,' *Slate* April 7, 2010. <http://www.slate.com/id/2248998/>
- DeGrazia, David, 'Moral Vegetarianism from a Very Broad Basis,' *Journal of Moral Philosophy* 6 (2009), pp. 143-165.
- , *Taking Animals Seriously* (Cambridge, UK: Cambridge UP, 1996).

³⁰ I am indebted to David Plunkett, Christian Coons, John Basl, Sean Walsh, Derek Baker, Mark Budolfson, Tom Dougherty, Gideon Rosen, Katie Batterman, two anonymous referees for this journal, an audience at Bowling Green State University, and the students in my Fall 2009 Ethics and Society class for helpful comments and discussion of ideas that went into this paper. Work on this paper was supported by a fellowship with the *Many Moral Rationalisms* research project, funded by the Australian Research Council.

- Descartes, Rene, *The Philosophical Writings of Descartes Vol. III*, ed. and trans. J. Cottingham, R. Stoothoff, D. Murdoch, and A. Kenny (Cambridge, Cambridge UP, 1991).
- Harman, Gilbert, *The Nature of Morality* (New York: Oxford UP, 1977).
- Herzog, Hal, *Some We Love, Some We Hate, Some We Eat* (New York: Harper Collins, 2010).
- Kahane, Guy, 'Evolutionary Debunking Arguments,' *Nous* 45.1 (March 2011), pp. 103-125.
- Kant, Immanuel, *Lectures in Ethics*, eds. P. Heath and J. Schneewind, trans. P. Heath (Cambridge: Cambridge UP, 1997).
- Korsgaard, Christine, 'Fellow Creatures,' *Tanner Lectures on Human Values* 25/26, ed. G. B. Peterson (Salt Lake City: University of Utah Press, 2004).
- , 'Realism and Constructivism in Twentieth-Century Moral Philosophy,' *Journal of Philosophical Research* APA Centennial Supp: Philosophy in America at the Turn of the Century (2003), pp. 99-122.
- Kutz, Christopher, *Complicity* (Cambridge: Cambridge UP, 2000).
- Lippert-Rasmussen, Kaspar, 'Two Puzzles for Deontologists,' *Journal of Ethics* 5.4, (2001), pp. 385-410.
- Marquis, Don, 'Why Abortion is Immoral,' *Journal of Philosophy* 86.4 (April 1989), pp. 183-202.
- McMahan, Jeff, 'Eating animals the nice way,' *Daedalus* 137.1 (Winter 2008), pp. 66-76.
- McPherson, Tristram, 'Moorean Arguments and Moral Revisionism,' *Journal of Ethics and Social Philosophy* 3.2 (June 2009), pp. 1-24.
- , 'Unifying Moral Methodology,' *Pacific Philosophical Quarterly* 93 (2012), pp. 523-549.
- Norcross, Alasdair, 'Puppies, Pigs, and People: Eating Meat and Marginal Cases,' *Philosophical Perspectives* 18 (Ethics) (2004), pp. 229-245.
- Orr, Deborah, 'Slave Chocolate?' *Forbes* April 24, 2006.
- Parfit, Derek, *Reasons and Persons*, (Oxford: Clarendon, 1984).
- Pollan, Michael, *The Omnivore's Dilemma* (New York: Penguin, 2006).
- Rachels, James, 'The Moral Argument for Vegetarianism,' in J. Rachels, *Can Ethics Provide Answers? And Other Essays in Moral Philosophy* (Lanham, MD: Rowman and Littlefield, 1997), pp. 99-107.
- Rachels, Stuart, 'Vegetarianism,' in T. Beauchamp and R. G. Frey (eds.) *The Oxford Handbook of Animal Ethics* (Oxford: Clarendon: 2011), pp. 877-905.
- Rawls, John, 'Outline of a Decision Procedure for Ethics,' *Philosophical Review* 60 (1951), pp. 177-97.
- Regan, Tom, 'The Case for Animal Rights,' in T. Regan and P. Singer (eds.), *Animal Rights and Human Obligations*, 2nd edn (Englewood Cliffs NJ: Prentice Hall, 1989), pp. 105-114.
- Ross, Jacob, 'Rejecting Ethical Deflationism,' *Ethics* 116 (July 2006), pp. 742-768.
- Scanlon, T. M., *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge, MA: Belknap, 2008).
- , *What We Owe to Each Other* (Cambridge, MA: Belknap, 1998).
- Sepielli, Andrew, 'What to Do When You Don't Know What to Do,' in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics* Vol. 4. (Oxford: Oxford UP, 2009), pp. 5-28.
- Singer, Peter, *Animal Liberation* (New York: Avon, 1977).
- , and Jim Mason, *The Ethics of What We Eat* (Rodale, 2007).
- Sinnott-Armstrong, Walter, *Moral Skepticisms* (New York: Oxford UP, 2006).

- Thomas, Janice, 'Does Descartes Deny Consciousness to Animals?' *Ratio* (new series) 19.9 (September 2006), pp. 336-363.
- Williams, Bernard, 'A critique of utilitarianism,' in J.J.C. Smart, and Bernard Williams, *Utilitarianism: for and against* (Cambridge: Cambridge UP: 1973), pp. 77-150.
- Wood, Allen, 'Kant on Duties Regarding Non-Rational Nature,' *Proceedings of the Aristotelian Society* Supp. Vol. LXXII (1998), pp. 189-210.