



MONASH University

A MEATY PROBLEM

How people think about the morality of meat

Lucy Mayne
B.IT., B.Litt., B.A. (Hons)

A thesis submitted for the degree of Master of Arts at
Monash University in 2018
Department of Philosophy

COPYRIGHT

© Lucy Mayne 2018

ABSTRACT

This thesis explores moral psychology, moral motivation, and the role of social norms in explaining both why meat eating is generally not recognized as a moral issue and why people tend not to be motivated to stop eating meat even if they become intellectually convinced they should. Analysing the moral psychology of meat eating reveals that a lack of motivation to stop eating meat is not just a failure of reason, affect, or motivation. Rather, it is at least partially a product of both empirical and normative expectations about the people with whom one interacts. The prevailing pro-meat norms within one's society make it difficult, though not impossible, for one to be motivated to treat eating meat as a moral issue.

Despite caring about animals, most people eat meat with a clear conscience. I argue that for people who care about animal suffering, the available justifications for eating meat fail; they either require one to reject the moral badness of causing animals to suffer, to reject that one should try to be a moral person, or to significantly restrict the range of meat considered morally permissible. Nevertheless, because such justifications for meat eating are widely accepted, they serve to mitigate the cognitive dissonance arising from eating meat. Strategically avoiding unpleasant information about meat also serves this goal. Importantly, both strategic ignorance and the employment of justifications that are inconsistent with one's other beliefs require widespread acceptance in order to be effective.

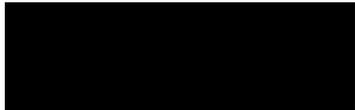
Many people explicitly acknowledge that their meat eating cannot be morally justified, yet eat meat with little or no guilt. The behaviour of such *conflicted carnivores* poses a particular puzzle. Since these people are motivated to act morally in many domains, an explanation from "weakness of will" is not compelling. Thus answering this puzzle requires explaining the different motivating force of different contextual domains. I argue, following Shaun Nichols, that moral motivation requires both a set of normative rules and an affective response. What, for the conflicted carnivore, distinguishes meat eating from other moral issues is the lack of affective fit between their moral beliefs about meat and those of their society.

Social factors play a key role in the practice of meat eating – eating meat is, for most people, not just an expression of individual preferences or a result of individual beliefs about the permissibility of the practice. Rather, descriptive norms govern what is available to eat and social norms make it socially difficult to avoid eating meat. Hence, if one is interested in changing meat-eating behaviour, it is not enough to persuade people that eating meat is immoral – these practical and

social risks also need to be addressed. Furthermore, because the types of meat which count as “food” are normatively defined, I argue that if moral vegetarianism were widespread then, over time, meat could be “taken off the table.” Thus in the right normative environment, eating meat could become moralized.

DECLARATION

This thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.



Lucy Mayne

April 3, 2018

ACKNOWLEDGEMENTS

This research was supported by an Australian Government Research Training Program (RTP) Scholarship.

Parts of this thesis were presented at the Australasian Association of Philosophy Conference 2016 and 2017, the Australasian Postgraduate Philosophy Conference 2016 and 2017, the Victorian Postgraduate Philosophy workshop 2016, the University of Melbourne Pre-AAP workshop 2017, the Monash Experimental philosophy group 2017, and milestone seminars at Monash University. I thank attendees of all of these events for their feedback, in particular, the panel members for my milestone seminars – Monima Chaddha, Bob Simpson, Linda Barclay, Catherine Mills and Graham Oppy. Thanks also to my peers Ben Groedeck, Maksymilian Sipowicz, and James McGuire, as well as my brother Tim Westcott for comments on drafts, and to Michael Smith and Shaun Nichols for helpful discussion.

A great thanks to my supervisor John Thrasher for detailed feedback and guidance throughout my thesis-writing process. Thanks also to my associate supervisor Rob Sparrow, for helping me to tighten the framing of the thesis.

For proofreading assistance, I thank Caitlyn Perry. I'm lucky to have a friend who can consistently pick up my split infinitives, even if I occasionally ignore her advice.

For keeping my sane throughout this mammoth undertaking, I thank the Melbourne (and Australia-wide) Modern Jive community for giving me a creative and social outlet to balance the long hours of reading and writing, with a special thanks to my wonderful dance partner Alysander Stanley. Special thanks also to Janet Coath for keeping me supplied with vegan cookies, intermittently throughout the thesis, and in double batches near the end.

Most importantly, thanks to my husband, Hunter Mayne. I don't know how I would have got through this undertaking without you.

CONTENTS

1	INTRODUCTION: MEAT EATING AS A PROBLEM FOR PHILOSOPHY	2
1.1	The puzzle of meat eating	2
1.2	Summary and preview	6
2	JUSTIFYING MEAT: THE PUZZLE OF GOOD-CONSCIENCE MEAT EATING	10
2.1	Why meat needs justification	10
2.2	Lay defences of meat	13
2.2.1	Direct justifications for meat – the 4Ns	13
2.2.1.1	Necessary	14
2.2.1.2	Natural	16
2.2.1.3	Normal	18
2.2.1.4	Nice	18
2.2.2	No-effect justifications	20
2.3	The inadequacy of philosophical justifications	22
2.3.1	Arguments that achieve too little	23
2.3.2	Arguments that achieve too much	26
2.4	Maintaining good-conscience meat eating	28
2.5	Conclusion	32
3	THE PUZZLE OF THE CONFLICTED CARNIVORE	35
3.1	Meat is bad, pass me the chicken	35
3.2	The conflicted carnivore	37
3.3	A puzzle on any view of morality	38
3.3.1	Internalist positions	39
3.3.2	Externalist positions	40
3.3.3	The common threads	41
3.4	Components of moralized judgement	42
3.4.1	Moralization	43
3.4.2	Reasoning against meat	45
3.4.3	The importance of affect	46
3.4.4	The problem with an affective explanation	49
3.4.5	Contextualising affect: the role of convention	51
3.5	Conclusion	55
4	NORMS OF MEAT	57
4.1	Meat eating as a collective behaviour	57
4.2	Meat as socially endorsed	58

4.3	Analysing the collective behaviour of meat eating	64
4.3.1	Types shared behaviours	65
4.3.1.1	Customs	65
4.3.1.2	Descriptive norms	65
4.3.1.3	Social Norms	66
4.3.2	Categorizing meat eating	67
4.4	Changing norms of meat	70
4.4.1	Unravelling pro-meat norms	70
4.4.2	Establishing anti-meat norms	74
4.5	Moralizing meat	75
4.5.1	The current moral status of meat	76
4.5.2	Changing the categorization of meat	78
4.6	Conclusion	80
5	CONCLUSION	83
	BIBLIOGRAPHY	89

'I just don't want to eat an animal that's standing there inviting me to,' said Arthur, 'It's heartless.'
'Better than eating an animal that doesn't want to be eaten,' said Zaphod.
'That's not the point,' Arthur protested. Then he thought about it for a moment.
'Alright,' he said, 'maybe it is the point. I don't care, I'm not going to think about it now. I'll just ... er ... I think I'll just have a green salad,' he muttered.

—Douglas Adams, *The Restaurant at the End of the Universe*

INTRODUCTION: MEAT EATING AS A PROBLEM FOR PHILOSOPHY

1.1 THE PUZZLE OF MEAT EATING

Most people around the world eat meat, or would if they could afford it. Most of those who eat meat do so with a clear conscience. Yet most people also care about animals, find animal suffering upsetting, and think it morally bad to cause animals to suffer, at least without justification. There is therefore a disconnect between how people view the moral significance of animal suffering, including that of animals killed for meat, and the moral complacency with which most people eat meat. A similar disconnect can also be seen regarding other widespread issues that arguably have significant moral implications, including climate change, extreme poverty, and the exploitation of workers in developing countries. Although these issues might intellectually seem like morally pressing ones, many people deny them to be moral issues. Furthermore, even when acknowledging them as moral issues, people rarely treat them as urgent. Most people go about their lives feeling moved neither to make any significant changes to their own behaviour nor to advocate for improved policies. In this thesis, I use the case study of meat eating to explore the puzzle of this disconnect. I argue that the prevailing norms in most societies make it difficult, though not impossible, for someone to treat eating meat as a deeply moral issue.

This is a thesis about moral psychology, moral motivation, and the role of social norms in intuitive morality. It is not a thesis on substantive normative ethics. I do not, therefore, argue for the immorality of eating meat. Ethical arguments for a moral obligation to abstain from eating meat have already been extensively covered from a variety of angles, including consequentialist/utilitarian (Singer 1975), rights-based (Regan 2004), Kantian (Korsgaard 2004), contractarian (Rowlands 2002; Talbert 2006), virtue ethical (Hursthouse 2006), feminist (Adams 2000), and framework independent (DeGrazia 2009; Engel 2000; McPherson 2016; Rachels 1997). While there are philosophical benefits to refining such arguments and developing new ones, that is not my project. Instead, I seek to understand why existing arguments seem to have so little practical effect on agents who care about behaving morally.

This lack of practical success could be due to shortcomings in these arguments or even due to the vegetarian position being fundamentally mistaken. However, while some meat eaters have come to a con-

sidered position regarding their meat eating, most have not. Many rely on transparently poor arguments yet are confident their meat eating is justified; many others go through life without ever seriously considering the vegetarian position, let alone engaging with what philosophers have to say. More interestingly, a non-trivial proportion of meat eaters have concluded that they should not eat meat, but do so anyway. All these observations make a *prima facie* case that the lack of success of arguments against meat eating cannot be explained by mere argumentative inadequacy. I therefore turn to a moral psychological approach to understanding the puzzle of meat eating.

Despite obvious moral implications relating to animal suffering and death, meat eating is largely treated as a moral non-issue. Most people, including most meat eaters, are distressed by animal suffering and consider mistreating animals to be immoral (Allen and Baines 2002; Plous 1993). For example, the idea of someone torturing an animal for pleasure is abhorrent to such people. Yet their meat eating seems, to them, perfectly justified, even though it is difficult to dispute that the level of animal suffering involved in mainstream animal agriculture is immense (for descriptions of standard practices, see Foer (2009), Joy (2010), and Mason and Singer (1990) and Singer (1975)). Knowledge of the animal suffering behind their food does not noticeably detract from most people's enjoyment of it, nor from their self-perception as moral beings. As Bramble and Fischer (2016, p. 1) observe, "Of course, we don't want [the animals we eat] to suffer unnecessarily. But judging by standard [...] farming practices, this desire doesn't run very deep." People do not wish suffering and death to be inflicted upon animals, but they do want cheap and plentiful meat and, apparently, see this as a reasonable trade-off. Put simply, people's negative attitudes towards suffering inflicted on animals seem to indicate that the suffering of animals for meat production should be treated as at least a potential moral issue – one deserving moral consideration even if not overriding moral action. It is not.

The problem of meat is that most people do not see it as a moral issue despite their other moral beliefs. For the most part, meat eating is treated not as something regrettable but all-things-considered justified. Rather, it is treated as falling altogether outside the scope of morality. I therefore begin with the observation that most people find the idea of deliberately hurting animals morally repugnant but also find the idea of eating meat to be perfectly fine, and explore why it is that the former does not commonly lead to a negation of the latter. I seek to understand, through analysing the case study of meat eating, how reason, judgement, affect, belief, and motivation connect to behavioural outcomes on a potentially moral issue. I also seek to understand how and to what extent one's social environment influences individual moral responses. Eating meat, particularly that of inten-

sively farmed animals, gives rise to a puzzle: how is it that people who care about animals can eat meat with such moral ease?

It might seem that there is a simple explanation for such behaviour regarding meat eating and other comparable issues – people are weak willed or just do not care. However, such answers are inadequate. While meat eating is not, for the most part, treated as a moral issue, most people strongly value doing the right thing in many contexts. While we can readily observe much immorality, we can also observe much moral behaviour, and not just from exemplary individuals. Ordinary people largely refrain from theft, murder, lying, and other widely condemned activities, even where a cost-benefit analysis would suggest wrongdoing to be the rational course of action. There are many moral issues where compliance with ‘the right thing’, even if imperfect, can be expected. In these cases, a disconnect between belief and action is much rarer and more remarkable, and weakness of will is a much lesser factor. Thus why people are poorly motivated to avoid meat *despite being motivated about other issues* needs explaining.

Furthermore, many other issues are or have been treated as deeply immoral in some societies without being treated as such in others. Today slavery is appalling to the average Westerner, but it hasn’t always been. Torturing animals for fun is largely unthinkable today, but cat burning was once an accepted form of entertainment. Eating whale meat is morally repugnant for many Australians, but this meat is considered a delicacy in Japan. Similarly, all of us think certain issues deeply important despite them not registering with others as even being moral issues. Some people think abortion a grave moral wrong, others think a grave wrong is done when women are denied this option. Some people find homosexuality morally abhorrent while others see acting on one’s sexual orientation as morally neutral. Some people think gun-carrying irresponsible and immoral, others are morally outraged at the idea that gun rights might be restricted. While people diverge on what they treat as serious moral issues, there is nevertheless widespread convergence on the idea that some issues are morally serious and deeply motivating. On some topics, a disconnect like that experienced by many meat eaters is unremarkable; on others, it would be seen as evidence of something being deeply wrong with that person. The difference between these types of issues needs explaining.

This discrepancy is not confined to meat eating. People frequently claim to think one thing about what morality requires while doing the opposite. They think climate change an important issue but drive instead of taking public transport, think extreme poverty terrible while spending money on frivolities, and think the exploitation of sweatshop workers inexcusable while buying clothing manufactured by them. Yet although meat eating is just one of many issues that tend to be non-motivating, it is worth exploring. First, like the issues men-

tioned above, meat-eating behaviour is interesting in its own right, especially to someone who agrees, at least intellectually, that it is a serious moral issue.

However, even if one isn't interested in meat eating for its own sake, it has advantages as a model for understanding the broader puzzle of how people form moral judgements and what motivates them to act on those judgements. While living consistently with one's beliefs is rarely easy, it is easier for someone to adopt a vegetarian diet than for the average Westerner to cut their carbon emissions to a level consistent with safe global per-person emissions, or for a person to follow an intellectual commitment to a movement such as effective altruism to its logical conclusion. Unlike adopting a carbon-neutral lifestyle or donating all unneeded money to charity, vegetarianism and even veganism are possible while maintaining a relatively normal life. Hence, of these complex cases of puzzling behaviour, meat eating is one of the simpler ones to understand. It is therefore a sensible starting point for understanding the factors involved in cases where people are either intellectually or practically resistant to changing their behaviour to be more in accordance with their apparent values. I hope that the understanding of the puzzles involved in meat eating achieved in my analysis can then be extended to understanding inaction and indifference on issues like climate change and extreme poverty.

This project is informed by empirical work, but I do not provide a merely sociological or psychological explanation of the puzzle of meat eating. Rather, I draw on empirical work to inform long-standing philosophical debates on moral motivation, and apply recent empirically informed philosophical work on social and descriptive norms to the specific case of meat eating. I seek, through the case study of meat eating, to provide an account of how reason, affect, and social factors (especially norms) contribute to something being treated as a deeply moral issue. My analysis of the problem of meat eating both draws on and has implications for a variety of philosophical sub-fields, including cognitive dissonance theory, moral motivation, and social norm theory. I argue that the justifications commonly employed in defence of meat eating are a mechanism for cognitive dissonance reduction rather than genuine justifications ([Chapter 2](#)), that both affect and convention play important roles in causing someone to treat something such as meat as a genuine moral issue ([Chapter 3](#)), and that meat eating has the potential to be treated as morally serious, but only if norms surrounding meat eating change in specific ways ([Chapter 4](#)). Thus the questions philosophers ask about moral motivation need to be reassessed and refocused away from a purely individualistic perspective towards instead understanding individuals as socially embedded and heavily influenced by what others around them do, yet still having independent agency. Importantly, this way of framing

the problem remains agent-centred – the social environment plays an important but not determinative role, which makes the theory testable.

It is worth noting a few limitations necessary to streamline my argument. First, for simplicity, I write only on meat eating, but much of my analysis could also be extended to the consumption of dairy, eggs, and other animal products for both food and non-food purposes. Similarly, while what I say about meat eating has implications for similar cases like those discussed above, I do not extend my discussion to those cases except to draw useful analogies; my thesis focuses squarely on meat-eating behaviour. Second, I conflate animal suffering and animal death. Some meat avoiders are more concerned with one, some with the other, and similarly some meat eaters will be more sensitive to considerations from one side or the other. However, given contemporary mainstream farming practices, the two run together in practice, and so separating out these motivations is unnecessary. Third, I only focus on moral reasons for avoiding meat. Some people are motivated to abstain from meat for health or environmental reasons or due to religious prohibitions. While exploring the interplay of these reasons would be interesting, I do not have the space to do so here. Finally, my analysis is very much focused on meat eaters in the developed world, and most of the empirical research I draw on is limited to Anglophone countries. While my claims may apply beyond these populations, further research would be necessary to establish this.

1.2 SUMMARY AND PREVIEW

I argue that the way people think about meat is, for the most part, incoherent. Most people care about animals and think deliberately inflicted animal suffering morally bad. They do not, however, see meat eating as morally bad, despite its causal connection to animal suffering. The justifications people use tend to be ones they would not accept in other domains, and many people cheerfully eat meat despite claiming to think animal suffering for meat production a serious moral issue. Analysing the moral psychology of meat eating reveals that a person's lack of motivation to stop eating meat is not just a failure of reason, affect, or will. Rather, it is at least in part a product of one's social environment, influenced by both *empirical expectations* (what one expects others will do), and *normative expectations* (what one expects others expect one to do). It is easier to treat a behaviour as deeply immoral if one expects that few people will do that thing. Furthermore, it is easier to avoid something if one thinks one will suffer sanctions of some sort for that behaviour, even if this is merely the disapproval of those whose opinions one cares about. Thus the inconsistencies in how people think about meat eating compared to

other more straightforwardly moral issues are primarily enabled by the collective acceptance of meat eating. The fact that meat eating is normal produces a firewall between intellectual recognition of the problem and the behaviour actually *feeling* wrong.

I start, in [Chapter 2](#) by examining reason – the justifications put forward by both philosophers and lay meat eaters. I argue that while these justifications need not fail universally, they cannot do the necessary work for any meat eater who at least sometimes eats the meat of intensively farmed animals, values being moral, and thinks that unnecessarily causing animals to suffer is morally bad. Hence these justifications cannot absolve the ordinary meat eater in the developed world unless he or she is willing to change other beliefs. All available justifications either treat meat as a special case, require that the animals used for meat be entirely denied moral status, require that one eat much less meat than most people consume, or require that the animals raised for meat have much better lives and deaths, on average, than they currently do. Thus neither philosophical nor lay defences can adequately justify meat eating for most meat eaters, given their other beliefs. Rather, reasons are being marshalled to bolster the pre-existing intuition that meat eating must be morally acceptable. I then argue that these justifications nevertheless seem adequate because meat eating is so common. The widespread nature of meat eating makes it easy for meat eaters, through strategies of reducing cognitive dissonance without resolving it, to hold mutually incompatible positions – people can strategically avoid thinking about the problems with their justifications, as well as strategically avoid gaining inconvenient information.

I then turn, in [Chapter 3](#), to an in-depth examination of the *conflicted carnivore*. Such people present a puzzle for moral philosophy: they are people who openly admit that eating meat is morally wrong yet continue to enjoy eating meat while experiencing only minimal psychological or moral discomfort. Such people reveal that it is not enough to rationally persuade someone that justifications such as those discussed in [Chapter 2](#) fail and that eating meat (or, indeed, any activity) is immoral – motivation will not necessarily follow from such a realisation. I argue that explanations such as insincerity of judgement or weakness of will are inadequate. Such explanations are proximate rather than ultimate; they fail to say anything interesting about what it is that distinguishes the conflicted carnivore's response to meat from either the committed vegetarian's or from their own response to issues they deem morally important. The conflicted carnivore has reasoned their way to a conclusion, but this alone is insufficient to motivate concordant action – he or she finds it easy to turn away from thoughts of animal suffering. I argue that this is because, while meat has been intellectually judged to be wrong, it does not *feel* wrong. Because of the importance of something feeling wrong,

an affective response is necessary to generate a strong moral judgement. However, affect is not sufficient for a moralized judgement because affect can be suppressed, as it often is with animal suffering. I then argue that the best way of making sense of the conflicted carnivore's disconnect between expressed moral judgement and behaviour is that social norms endorsing meat eating encourage the suppression of negative affective responses to meat such that these affective responses are not seen as indicative of something being *wrong*.

Finally, in [Chapter 4](#), I analyse the behavioural regularity of meat eating. I argue that eating meat is not simply a means of satisfying an individual desire, but is instead heavily influenced by the meat-eating practices of others. Because of the operation of norms, giving up meat has many elements of a collective action problem. Early adopters of vegetarianism face not only practical problems, but also the risk of social sanctions, making acting early unattractive. Thus the lack of moral motivation observed in many meat eaters is not straightforwardly an individual failure, but rather derives from the structure of norms and expectations extant in that person's society. As long as meat eating is widespread, it will not invite much, if any, disesteem, even from individuals who personally disapprove. Thus the reason people are not open to arguments about the immorality of a practice such as meat eating, and why they so often fail to act even when they apparently are persuaded, has more to do with social factors than is often acknowledged: a person's social environment is more important than individual judgements for predicting behaviour. This means, however, that if there is consensus within one's reference group, meat eating can be "taken off the table" such that the costs and benefits of meat eating cease to be considered. I argue that if moral vegetarianism were widespread, eating meat would invite disesteem, thereby motivating compliance, which could, over time, lead to the widespread moralization of meat.

Examining how people reason about meat eating reveals that pointing at individual explanations for moral failure, at least for wide-reaching issues, is inadequate. While people may use poor justifications for why they take part in an activity that might seem wrong given their other beliefs, their doing so is enabled by the widespread endorsement of such justifications, which in turn makes such behaviour seem more excusable. Furthermore, collective inaction makes it difficult to treat an issue as being deeply immoral even if one is intellectually persuaded that one should – moralization requires an affective response to the breaking of an endorsed rule, which is difficult if social conditions both encourage desensitization and explicitly allow that behaviour. Hence understanding the norms surrounding a practice is essential to understanding why it is or is not moralized. Furthermore, if one's goal is not just understanding but also effecting change, then this is a necessary first step. Philosophers' analyses of

moral 'failings' therefore need to take into account the social framework within which individuals operate.

JUSTIFYING MEAT: THE PUZZLE OF GOOD-CONSCIENCE MEAT EATING

2.1 WHY MEAT NEEDS JUSTIFICATION

The moral implications of eating meat are generally far from people's minds as they sit down to dinner. Yet most people hold beliefs that are, *prima facie*, in tension with moral acceptance of their own meat eating. Nevertheless, most meat eaters are comfortable with this behaviour – 89% of respondents in a recent study I conducted with John Thrasher agreed that meat eating is morally acceptable (Thrasher and Mayne 2018).¹ In this chapter, I argue that for a significant proportion of meat eaters, their belief that it is morally acceptable for them to eat intensively farmed meat cannot be reconciled with their other beliefs about morality and the moral status of animals. Such meat eaters are in a state of cognitive dissonance with respect to their meat eating – their beliefs about the moral status of animals cannot be straightforwardly reconciled with their belief in the moral permissibility of meat eating. I then argue that although the justifications commonly employed by meat eaters fail to establish the moral permissibility of meat eating given other beliefs held, they nevertheless serve to alleviate cognitive dissonance. Furthermore, the widespread acceptance of meat eating makes it seem excusable, and is a key factor in making the available justifications succeed in alleviating cognitive dissonance.

Meat eaters tend to hold beliefs that are *prima facie* at odds with them seeing meat eating, as it is currently practised, as being morally acceptable. Many meat eaters see themselves as animal lovers, and have close relationships with their pets or a fondness for animals more broadly. Even among those who dislike animals, few endorse animal suffering, at least in the abstract – 97% of our study participants agreed that other things being equal, animal suffering is bad (Thrasher and Mayne 2018). Yet 86% of our survey respondents thought both that eating meat is morally acceptable and that animal suffering is bad (Thrasher and Mayne 2018). Meat eaters, particularly the majority who eat meat from intensively farmed animals, “engage in a diet that requires [animals] to be killed and, usually, to suffer” (Loughnan, Bastian, and Haslam 2014, p. 104). Thus the animal suffering involved in meat production should be a hurdle for meat eaters who care about animals.

¹ Participants ($n=863$) were recruited via Amazon Mechanical Turk, and were mostly from the US. Data available on request.

Because of this, meat eaters who care about animals may sometimes, when confronted with the suffering caused by their diet, feel uncomfortable about its moral implications, experiencing a tension between their enjoyment of meat and their self-image as loving and compassionate people who are against animal cruelty. This tension can lead to ambivalence: a study by Berndsen and Van Der Pligt (2004, p. 71) found that 69% of meat eaters simultaneously experience both positive and negative feelings toward meat or meat eating. Meat is seen as good and tasty, but meat involves animal suffering, and deliberately inflicted animal suffering is seen as morally bad. Thus a significant proportion of meat eaters agree that it would be morally better if they ate less meat (45% of respondents in our survey) and that it would be morally better if most people ate less meat (50% in our survey) (Thrasher and Mayne 2018). Similarly, 49% of US adults support a ban on factory farming and 33% support a ban on all animal farming, despite 97% agreeing that “Whether to eat animals or be vegetarian is a personal choice, and nobody has the right to tell me which one they think I should do” (Reese 2017). These attitudes appear inconsistent: a non-trivial proportion of meat eaters believe their meat eating to be morally justified, or at least permissible, while holding beliefs incompatible with this position.

Of course, something can be bad in isolation while being all-things-considered justified. Thus, while meat eaters see a problem with animal suffering, they consider eating meat to be morally permissible; indeed, they must do so to eat it in good conscience. In this chapter, I explore how this is managed. I first show that the most common lay justifications for meat eating should not be able to make meat eating seem all-things-considered justified for a significant proportion of meat eaters. I then show that the best available philosophical arguments for meat eating also cannot, for most meat eaters, justify meat eating as it is currently practised. My aim here is not to show that all arguments in favour of meat eating fail. Rather, I show that most people’s meat eating cannot be justified, *given their other beliefs*. Specifically, I argue that meat eating cannot be justified for anyone who holds the following three beliefs, which together form what I call the *meat trilemma*:²

1. *ceteris paribus*, causing animals to suffer is morally bad because of what it does to the animal – animals have moral status;
2. *ceteris paribus*, one should try to be a moral person and not act in ways that have morally bad results; and
3. eating meat, including the meat of animals who were raised on factory farms or can otherwise be reasonably inferred to have led bad lives, is morally permissible.

² In this, I loosely follow Milan Engel’s (2000) argument for vegetarianism from beliefs most people already hold.

Although the meat trilemma does not apply to all meat eaters – any or all of these beliefs can be rejected – given our survey results, this trilemma likely applies to a significant proportion of meat eaters. The meat eating of people caught in the meat trilemma should be puzzling.

Given the third belief of the meat trilemma, it is also plausible that moral issues with meat eating might be eliminated through lab-grown or ‘humanely raised’ meat. This would, in theory, enable meat judgement 3 to be replaced with:

3. Eating meat is morally permissible where producing this meat has not caused any significant animal suffering.

However, while these options may offer possible solutions, they are not currently available to the ordinary meat eater. Lab-grown meat is not yet commercially available. As for ‘humane’ meat, there is a question of whether humane guidelines are sufficient to ensure an adequate quality of life,³ and even if they are, the moral permissibility of ‘killing happy animals’ still requires argument, something beyond the scope of this thesis. Nevertheless, for the purpose of this chapter, it is enough to note that even if ‘humane’ standards are thought to be adequate, very few people consistently eat meat *only* from humane sources. Thus insofar as someone at least sometimes eats meat from non-humane sources, eating that meat poses a problem for people who endorse beliefs 1 and 2 of the meat trilemma.

Nevertheless, to most meat eaters, at least some of the justifications for meat eating seem obviously right. However, I argue that the best available justifications, both lay and philosophical, fail to resolve the meat trilemma. They all either rely on false premises or require the meat eater to give up or significantly and unpalatably alter either belief 1 of the meat trilemma, and downgrade the moral status of animals, or belief 3, and heavily restrict the meat they consider permissible to eat. Thus for someone caught in the meat trilemma, typical meat-eating behaviour cannot be justified.

Consideration of the lay arguments, however, reveals that even though they fail to resolve the meat trilemma, these justifications help alleviate cognitive dissonance. Cognitive dissonance is the unpleasant feeling one experiences when one does something inconsistent with one’s beliefs (Rabin 1994, p. 178). Even though the lay arguments cannot clearly justify meat eating, they can ward off any negative feelings by making meat eating *seem* justified, or at least excusable, especially as there is moral wriggle room between something being morally *bad*, and it being morally *wrong* such that one would be a bad person for

³ Even certified humane meat is subject to production pressures which result in the need to trade off animal welfare against profitability. For example, the most stringent humane certification standard for meat chickens in Australia, the ACO, allows a stocking density of 12.5 birds per square meter, or 80cm² per bird (Herron 2015), still a crowded existence.

doing it.⁴ Employment of strategic ignorance can further improve cognitive dissonance management – meat eaters can deliberately avoid learning about the lives of the animals they eat, and can also avoid information that might undermine their pro-meat justifications. Importantly for my broader argument, both these approaches require a critical mass to be effective. If most people rejected the justifications for meat eating, it would be difficult for one to endorse them; if few people ate meat, it would be difficult to excuse it; if most people were informed, then ignorance would be much harder to maintain.

2.2 LAY DEFENCES OF MEAT

Most people who eat meat will, if questioned, present justifications for their meat eating. In this section, I canvass the most common lay justifications for meat eating and show they cannot resolve the meat trilemma. This is a negative project. Rather than demonstrating that justifying meat eating is impossible, I merely aim to show that the justifications most commonly put forward by everyday meat eaters fail: they are either easily refutable on factual grounds (in the case of strong necessity) or they require giving up beliefs constitutive of the meat trilemma in ways unacceptable to most meat eaters. Specifically, since the meat eater who invokes these justifications is operating within a framework of assumed morality, all examined lay justifications require either the first belief of the meat trilemma to be given up, or the scope of permissible meat to be significantly reduced. In [Section 2.3](#), I then argue that a similar problem exists for the more philosophically sophisticated arguments defending meat.

2.2.1 *Direct justifications for meat – the 4Ns*

Recent psychological research reveals that most explicitly stated justifications for eating meat fall into one of four categories: *Normal*, *Natural*, *Necessary*, and *Nice*. 83-91% of the top 3 spontaneously offered justifications for meat eating fall into one of these 4 categories (Piazza *et al.* 2015), which together form the 4Ns of meat. Insofar as they are taken to be valid, these justifications enable an outright rejection of the immorality of meat. However, none of the 4Ns can block the vegetarian conclusion without either relying on easily refutable premises or on assumptions denied by those same people in other contexts. Use of the 4Ns therefore requires either a wholesale rejection of morality (a rejection of belief 2 of the meat trilemma), or a denial that animals matter morally and hence that causing them to suffer is bad (a rejection of belief 1). Hence, both individually and jointly, justifications in the 4N categories fail to provide a resolution to the meat trilemma.

⁴ In [Chapter 4](#), I directly examine the role social norms play in shaping what is seen as wrong, and what seems merely bad.

2.2.1.1 *Necessary*

Arguments grounded in necessity would offer the strongest justification for meat eating, but they rest either on false premises or shaky inferences. Necessity, *prima facie*, offers the strongest justification for any activity. Necessity-grounded justifications for meat eating include “It is necessary to eat meat in order to be healthy”, “A healthy diet requires at least some meat”, “You cannot get all the protein, vitamins and minerals you need on an all plant-based diet”, and “Human beings need to eat meat” (Piazza *et al.* 2015, p. 118). High meat consumption is correlated with strong beliefs that meat eating is necessary for health (Rothgerber 2012). Whether belief in necessity causes heavy meat eating or heavy meat eating encourages endorsement of necessity, the necessity justification clearly plays an important role.

If meat consumption were necessary for survival (or for reasonable health) then meat eating would be inevitable; vegetarianism would be tantamount to a slow form of suicide (Joy 2010, p. 109). Hence strong necessity could resolve the meat trilemma through introducing a higher-order consideration, a justification overriding the *ceteris paribus* part of the first belief of the meat trilemma. Assuming for the sake of the argument the truth of Kant’s maxim that ‘ought implies can’, the necessity of meat would mean that refraining from eating meat would, at most, be supererogatory in the most heroic and self-sacrificial of ways; it could not be morally required. Thus participation in a practice that harmed others who have moral standing could be morally justified.

However, the strong necessity claim – that eating meat is necessary for survival or decent health – is false. While planning a balanced vegetarian diet may be more difficult and require more planning than eating meat, especially in a culture where this is not the norm, there is good reason to believe that healthy vegetarianism is, at least in theory, accessible to most people. For example, the American Dietetic Association states that:

appropriately planned vegetarian, including vegan, diets are healthful, nutritionally adequate, and may provide health benefits in the prevention and treatment of certain diseases. These diets are appropriate for all stages of the life cycle, including pregnancy, lactation, infancy, childhood, adolescence, older adulthood, and for athletes. (American Dietetic Association 2016, p. 1266)

Similarly, while the Dietitians Association of Australia emphasises the need for planning, they state that “A vegetarian diet can be healthy as many plant foods are low in saturated fat and high in dietary fibre” (Dietitians Association of Australia 2016b) and that “it is still possible to obtain all the nutrients required for good health on a vegan diet” (Dietitians Association of Australia 2016a). Hence the best

dietary evidence currently available indicates that eating meat is not necessary for health for most people, and indeed, that vegetarianism offers health benefits. Even for someone unwilling to conduct such research, the strong claim that meat is necessary for survival in decent health is also readily refuted through a visible counter example: healthy long-term vegetarians and vegans. Meat is clearly not universally necessary for health or survival, and thus strong necessity is not available to most meat eaters.

Most of the time, however, necessity is not being employed in the strong sense. Rather, “necessary” is being used to make the weaker claim of “easier”. The claim is something along the lines of “Given my goals and the amount of effort I am prepared to put into achieving them, it is necessary for me to eat meat.” Like strong necessity, weak necessity aims to show that the harm to animals for meat production is a justified bad rather than a moral wrong. For example, someone unwilling to put time into researching plant-based nutrition might claim that meat eating is necessary for them to achieve the goal of remaining healthy. However, this invocation of weak necessity presupposes not only that the meat eater’s goals are sufficiently important to justify the animal suffering involved, but also that the effort differential between the chosen path and a less harmful one is sufficient to justify the animal suffering involved. In many cases, as with the meat eater who desires to be healthy, it is plausible that the importance of one’s goals could justify causing harm, absent alternatives. However, the effort differential between the harm-causing option and the next-best option remains morally relevant.

If one can easily achieve the same or a similar goal while doing less harm, it is hard to justify the harm as “necessary”, even under weak necessity. It is not clear that the effort involved in researching meat alternatives is high enough to do the justificatory work required by belief 1 of the meat trilemma – if the effort involved is thought to justify meat eating, then one must assign a very low importance to the suffering of animals for meat production. Yet in other contexts, the same people assign a much higher importance to such animal suffering. For example, many meat eaters refuse to wear fur for animal welfare reasons. Nevertheless, it is noteworthy that the amount of effort involved in pursuing a less harmful path will be heavily influenced by how prevalent alternative behaviours are, something I return to in [Chapter 4](#). Furthermore, even if strict vegetarianism is too effortful to be expected of someone, most meat eaters could eat significantly less meat (for example, choosing a vegetarian option where it is available) without expending any additional effort, making it implausible that the effort involved (and hence weak necessity) can resolve the meat trilemma.

More interestingly, there is an asymmetry between how the necessity claim is deployed to defend meat eating and how necessity func-

tions in other moral contexts. Common sense morality tells us that necessity can sometimes override the wrongness of an act. For example, most people agree that necessity would justify killing one's assailant in self-defence, and this agreement is reflected in criminal codes which excuse someone who kills in self defence (e.g. Criminal Code Act (1995, §10.4)). More reluctantly, many would concede that necessity would justify smothering a baby to try to save the lives of a group of people who would all be killed if discovered, as is heart-wrenchingly described in Aaron Elster's memoir of his childhood experience of the holocaust (Elster and Miller 2008). However, in the staple moral tragedies of philosophical discussion, necessity is invoked to justify an action which most of us would be reluctant to carry out. Whatever the extenuating circumstances, few people desire to smother a baby or to kill in self-defence, and most of us would struggle psychologically for the rest of our lives with having been forced to act in such a way. We would wonder if our action had truly been necessary or if some alternative might have existed. Meat, however, is different. In the case of meat eating, necessity is invoked to justify *something we wish to do anyway*. The necessity of meat is not seen as regrettable; it is enthusiastically embraced, and its invocation requires assigning a very low moral status to the animals used for food, and hence a (near complete) rejection of the idea that animals matter morally. Necessity does not, therefore, resolve the meat trilemma.

2.2.1.2 *Natural*

Unlike arguments from necessity, justifications relying on the naturalness of meat eating need not rest on false beliefs. Justifications from naturalness take forms along the lines of "It is only natural to eat meat", "Our human ancestors ate meat all the time", "It is unnatural to eat an all plant-based diet", and "Human beings are natural meat-eaters – we naturally crave meat" (Piazza *et al.* 2015, p. 118). Such statements may be true, but this in itself says nothing about the morality of meat eating. Naturalness fails to bridge the is-ought gap. Thus unlike strong necessity, it does not provide a sufficiently strong justification to override belief 1 of the meat trilemma.

Naturalness is treated as offering strong support for meat eating, yet whatever support naturalness offers can be overridden by other considerations. Harm (e.g. to a raped woman, despite the 'naturalness' of a man's sex drive) and social undesirability (e.g. public nudity, despite clothing being 'unnatural') are both taken to make natural things wrong or undesirable. Hence for naturalness to justify a practice, there needs to be a principled way of differentiating between something natural and therefore good (as meat eating is alleged to be) and cases where something natural is nevertheless bad (such as rape

or public nudity). It is this differentiating mechanism, rather than naturalness per se, which is doing the justificatory work.

It is doubtful that naturalness adds any weight to the permissibility of a practice. Most writers on liberty have endorsed a “presumption in favour of liberty”: individuals should be left free to make their own choices, other things being equal (Feinberg 1984, p. 9). Thus if something is not harmful, it should be permissible and naturalness is not needed to further support it. Conversely, if something is harmful, naturalness alone cannot counteract this, as the above examples show. Hence justifications from naturalness can only be thought to work if the harm caused to the animals eaten is assumed to be outside the moral calculus. This assumption, however, is at odds with the care for animals exhibited by those caught in the meat trilemma – it requires a rejection of the first belief. Thus naturalness, as it tends to be employed, cannot enable an escape from the meat trilemma.

Closer examination of natural-type justifications, however, reveals they are less about what is *natural* and more about what *seems normal* and is therefore thought to be good. What is natural and what is normal are conflated, and it is the normalcy of a practice that differentiates natural-and-good from natural-but-bad. An example of this conflation is homosexual activity. Homosexuality is frequently referred to as unnatural by its opponents, yet same-sex copulation is observed in the animal kingdom (Bagemihl 1999). Homosexuality is therefore not unnatural in the sense of “not found in nature”. Rather, it is against the norms seen as important by those who oppose it.⁵ In the case of meat, it is true that nearly all pre-modern humans ate meat, making it natural in that sense. However, meat eating does not seem natural to people in non-meat eating cultures (for example, 55% of Indian Brahmins are lacto-vegetarian (Yadav and Kumar 2006), with vegetarianism being dominant in some cultural groups in certain regions of India). Nor does it seem natural to many second generation vegetarians in Western societies. Similarly, while it is common to claim that it is natural to want to eat meat, few Australians think it is natural to want to eat horse-meat or dog-meat, even though eating these meats is common in other cultures.⁶ It is the normalcy of meat eating in one’s culture, rather than what would happen in a state of nature, that gives rise to a sense of naturalness, as I will further investigate in [Chapter 4](#).

⁵ This is compatible with usage in natural law theory, which uses ‘natural’ not to mean “found in nature” but rather something along the lines of “in accordance with the laws of reason” (Finnis 2015, §1.1.1).

⁶ Hence, in order to avoid offending Western tourists, both Chinese and South Korean officials took steps to reduce the visibility of dog meat during, respectively, the 2008 Olympics in Beijing and the 2018 Winter Olympics in PyeongChang (BBC 2008; McCurry 2017).

2.2.1.3 *Normal*

The *normal* set of justifications rests on the idea that what is commonly done must be morally acceptable or at least not seriously wrong: either the majority cannot be wrong, or even if they are, following along with them is excusable. Reasons proffered in this category take forms along the lines of “It is normal to eat meat”, “It is abnormal for humans not to eat meat”, “Most people eat meat, and most people can’t be wrong”, and “It is common for people to eat meat in our society, so not eating meat is socially offensive” (Piazza *et al.* 2015, p. 118). Normalcy provides a cover – if it is commonplace to do something wrong, then it is easy to believe that behaviour is not seriously wrong and that one is not a bad person for engaging in it.

Arguments from normalcy are applied in an ad hoc manner to support things that receive social approval while not being seen as excusing ‘normal’ things that do not. Many practices now thought to be highly morally objectionable were once ‘normal’, including slavery, racial apartheid, the subordination of women, foot-binding, and the beating of children and animals (especially beasts of burden). However, most people who defend meat eating by appeal to normalcy would not agree with the relativistic claim that slavery, foot-binding, child marriage and other such practices were morally acceptable when and where they were the norm. If an appeal to normalcy would not persuade such people of the moral acceptability of those practices, then normalcy itself is not doing the justificatory work. Rather, familiarity with a practice makes it seem morally acceptable, such that *our* normal seems justified while *their* normal does not. Insofar as a practice is seen as immoral, normalcy can serve at most as an excuse, but not as a moral justification; an appeal to normalcy thus presupposes rather than establishes the moral acceptability of a practice. Justifying meat eating through normalcy presupposes that the harm caused to animals in meat production is not a moral issue. Yet this assumption is in tension with the first belief of the meat trilemma. Therefore normal-type justifications cannot enable an escape from the meat trilemma, though they provide an insight into the psychology of meat eating, one that will be further explored in Chapters 3 and 4.

2.2.1.4 *Nice*

The pleasure gained from meat factors strongly in the justifications given for meat eating. Indeed, enjoyment of meat had the highest level of agreement in a study of Australian motivations for meat eating (Lea and Worsley 2003, p. 508). Reasons put forward in the *nice* category take forms along the lines of “Meat is delicious”, “Meat adds so much flavor to a meal it does not make sense to leave it out”, “The best tasting food is normally a meat-based dish (e.g., steak, chicken

breast, grilled fish)", and "Meals without meat would just be bland and boring" (Piazza *et al.* 2015, p. 118). It is strange, however, for niceness to be acting as a moral justifications at all.

Pleasure is important. Intuitively, it needs to be allowed for in a plausible moral framework. Nevertheless, the pleasure gained from eating meat cannot do sufficient justificatory work to establish its permissibility. It is widely accepted, including by most meat eaters, that causing definite harm to others in order to gain pleasure is not morally acceptable. For example, sadism is seen as morally problematic because it involves non-consensually hurting others who matter morally as a means to gaining pleasure. Furthermore, it seems problematic even if the sadist's pleasure exceeds the harm to the victim. Thus arguing that eating meat is acceptable or even good because of the pleasure it brings presupposes that the suffering of animals either does not matter at all, or matters extremely little.

Furthermore, while eating meat can give gustatory pleasure, the relevant pleasure differential is not between eating meat and eating nothing, but rather between eating meat and eating a similar vegetarian dish. Thus even if one holds that vegetarian meals cannot be as good as those with meat, the differential is much less than commonly asserted (Engel 2000, p. 878). Given this, for anyone able to enjoy vegetarian food, animal suffering would have to be considered of extremely low or no importance for niceness to justify meat consumption. Yet this is at odds with the moral value assigned to animals in the first belief of the meat trilemma.

Many meat eaters would be horrified at the suggestion that pleasure would justify directly mistreating an animal. Someone who gains enjoyment from inflicting suffering on animals (including animals of species used for food, let alone dogs or cats) in a way not sanctioned by society will be seen by most as morally bankrupt.⁷ Yet many of these same people use pleasure to justify eating similarly mistreated intensively farmed animals. In order to succeed, nice-type justifications therefore require the animals used for meat to be assigned a lower moral status than they are in other contexts. Niceness therefore cannot justify meat eating unless one is prepared to concede that hurting animals for pleasure is justified for sport, entertainment, and sadistic pleasure, and not just for food. As with justifications in the normal and natural categories, the invocation of nice-type justifications presupposes the moral acceptability of meat eating in order to work, and so cannot resolve the meat trilemma –

⁷ Pleasure is, however, a component in justifying other forms of suffering inflicted on animals where it is sanctioned within one's reference group. See, for example, Roger Scruton's (1998, Appendix 2) defence of fox hunting, or Hal Herzog's (2010, p. 149-173) discussion of cockfighting.

the only way it can work is if belief 1 is given up, a difficult bullet to bite.⁸

2.2.2 *No-effect justifications*

Above, I examined the most common approaches to justifying meat eating directly. However, there is another common form of justification, the *no-effect* justification, which accepts that meat production is morally bad (or even wrong) but denies a personal moral obligation to refrain from eating meat. Justifications that support this distinction question the efficacy of personal abstention, claiming that the dietary habits of one person (or one family) contribute little, if anything, to reducing animal suffering in the meat industry. Such justifications serve to block the connection between the purportedly negative act (eating meat) and the negative outcome (animal suffering) – they deny the implicit premise in the meat trilemma that connects the first and third beliefs. The animal suffering involved in meat production can therefore be acknowledged to be bad without causal responsibility attaching to the meat eater. Without the meat eater being causally responsible for the harms involved, it might seem as though the meat trilemma has been dissolved. However, resolving the meat trilemma in this way neglects other important moral concepts and so involves inconsistent moral reasoning for anyone who is not a strict consequentialist.

The no-effect claim can take a variety of forms, ranging from the idea that once an animal has been killed, eating it cannot harm it, to the claim is that any individual changes in meat purchase and consumption will, due to noise in the supply chain from farmer to abattoir to supermarket to consumer, have no effect whatsoever on the number of animals raised and killed for meat (e.g. Budolfson (2016)). Even if various no-effect claims are empirically true, the way they tend to be employed is inconsistent with much ordinary moral thinking; no-effect reasoning neglects moral concepts such as complicity and integrity. Thus most employers of the no-effect claim use it in a circular way. In assuming these concepts to be irrelevant in the case of meat eating, they assume that meat eating is permissible in order to justify meat's permissibility. Thus irrespective of the empirical truth of the no-effect claim, it cannot be straightforwardly used to justify eating meat for meat eaters who value these other moral concepts.

No-effect reasoning is entirely consequentialist: the claim is that since one's actions have little or no impact, one might as well continue. Yet few people approach morality from a purely consequentialist perspective. Folk morality places weight on various non-consequence-oriented intuitions; someone who considers only

⁸ Section 2.3.1 addresses a more philosophically sophisticated version of the Nice justification.

consequences is aberrant. Consider the following vignette borrowed from Julia Driver:

Blake believes that sentient animals have moral standing and that, if eating meat harms animals, one ought not to eat meat. However, she has recently read articles that indicate that her individual choice to eat meat on any given occasion makes no difference—it in no way reduces animal suffering. She also believes that if she were to eat meat, she would be part of something that does lead to devastating harms to animals, however, since her individual choice makes no difference, she decides to eat meat. Also, precisely because she does care about reducing animal suffering she makes sure to purchase meat only from large industrial factory farms, rather than small, local, more humane producers. She does this because she believes that her decision to buy meat from a small producer is more likely to have an impact on production practices. She believes that she will be the cause of suffering if she buys more humanely produced meat. Therefore, she should not buy humanely produced meat, but, rather, buy meat produced on a massive industrial scale in which huge suffering is the norm. But, as long as she causes none of it, that is perfectly okay. (Driver 2016, p. 68)

Blake is following impeccable no-effect reasoning, yet intuitively, there seems something unsavoury about her behaviour. It is reasonable to question Blake's sincerity (or that of someone like her) in her judgement that one ought not eat meat if doing so harms animals. It seems not a justification, but rather a rationalization for a desired behaviour (Driver 2016, p. 74).

People who utilize the no-effect justification reap the benefits of a practice they claim to condemn. In doing so, they neglect many considerations that would ordinarily factor into moral reasoning. If someone is using the no-effect claim sincerely, they admit that it would be a good thing if most people stopped eating meat. Hence their continued cheerful meat consumption shows "an unwillingness to engage in the cooperative enterprise of ending animal suffering" (Driver 2016, p. 74). Meat eating is not unique in this regard: the same unwillingness to act is a common response to other wide-ranging issues such as reducing carbon emissions or avoiding purchasing products made by exploited workers. This may be because it make sense to avoid being the schmuck who does good while nobody else does, something I return to in [Chapter 4](#). Nevertheless, the cheerful continuation of a practice these meat eaters think would be best abandoned is hypocritical at best.

Furthermore, there are other more straightforward violations of folk morality embedded in the no-effect justification. Meat eaters who

admit that eating meat would be wrong absent the no-effect factor are complicit in wrongdoing – they tolerate and benefit from a practice that is by their own admission wrongful (Driver 2016, p. 75), and wrongful such that they could not obtain what they want (cheap, readily obtainable meat) without that practice. Intuitively, complicity in wrongdoing seems to matter even absent any causal effect. Consider, for example, someone who ineffectively participates in a murder plot – the plot would have been just as successful without them. In such a case it nevertheless seems appropriate to hold that person blameworthy in some sense, even if causal responsibility is not rightly attributable (Driver 2016, p. 71). Furthermore, as Mark Budolfson (2016) notes, given current technologies, killing animals is *highly essential* to meat eating – one could not enjoy meat consumption with any regularity unless animals were being killed, and one could not enjoy cheap meat unless animals were subjected to harmful intensive agricultural practices. This makes meat different to clothing or electricity, where the harms involved are not essential to production.⁹ There also seems to be a lack of integrity displayed by a no-effect endorser – such a person follows the crowd because it is easy rather than doing what seems to them to be right. Hence attributing moral failing does not require also attributing causal responsibility.

The preceding considerations do not matter from a strictly consequentialist perspective, but few people make moral decisions on purely consequentialist grounds. Insofar as people employ non-consequentialist considerations such as complicity or personal integrity in other cases of moral reasoning, the exclusion of these considerations when applied to eating animals is ad hoc. Thus the no-effect justification can only rescue the strictly consequentialist meat eater. For someone who thinks, as most meat eaters do, that non-consequentialist moral considerations sometimes matter, taking the position that the causal ineffectiveness of avoiding meat oneself is sufficient to justify meat eating presumes that meat eating is a non-moral issue in order to establish its moral permissibility. Hence this approach also fails to resolve the meat trilemma with integrity.

2.3 THE INADEQUACY OF PHILOSOPHICAL JUSTIFICATIONS

Philosophers' arguments do not fare much better than those of laypeople; they too cannot provide a satisfactory escape from the meat trilemma. Many philosophers' arguments do not defend meat eating as it is practised, but rather defend the possibility of ethical meat consumption or the consumption of life-enriching meals. Such

⁹ As Budolfson (2016) notes, the distinction between complicity with essential harm and mere complicity with harm seems to capture how vegetarians who consume non-meat animal products justify doing so despite the significant levels of harm involved, particularly with caged eggs – unlike with meat, the harm is seen as non-essential.

arguments do not help the meat eater chowing down on meat not endorsed by the argument. Those arguments that defend status quo meat eating require the meat eater to accept unpalatable conclusions – they require a rejection of the idea that it can ever be wrong to cause animals to suffer *because of what it does to the animal*, and hence require a rejection of the first belief of the meat trilemma. Thus even if the arguments canvassed below succeed in defending meat eating, they do not help the average person caught in the meat trilemma.

2.3.1 *Arguments that achieve too little*

The most common pro-meat position in the philosophical literature is avocation for or defence of some form of conscientious carnivorism such as that argued for by Michael Pollan (2006), Terence Cuneo (2016) or Benjamin Lipscomb (2016). Although there are subtle differences in their positions, all argue that it is not in principle wrong to kill and eat animals: eating the meat of animals who have had a good life and a swift and painless death is morally permissible. Yet such a position does not provide a defence of meat eating as it is currently practised. Rather, conscientious carnivorous requires a rejection of mainstream farming practices. For example, Pollan writes:

Even if animals can't suffer like human beings, there is no excuse for the cruelty that goes on in our factory farms and feedlots. Believe me, the people who run those places don't waste any time thinking about animal suffering. If they did, they'd have to go out of business. (Pollan 2009, p. 255)

Thus even if such approaches succeed in establishing the moral permissibility of some meat eating, they do not provide a defence of meat eating as it is currently practised. Rather, they argue against it. Yet few people eat *only* meat that would meet these criteria. Thus few meat eaters can employ such an argument with integrity.

A similar line of argument is that raising animals for meat, which necessarily involves killing them, gives a chance at life to animals who would not otherwise have been born. For example, Christopher Belshaw (2016) argues that meat eating can be good for animals insofar as the practice causes animals to be brought into existence who have, on balance, good lives. This argument, too, requires that the animals in question live good lives. However, on this basis, he rules out mainstream farming practices:

Animals in factory farms suffer a premature death, considerable pain throughout their lives, and considerable, and discomfiting, restrictions on their freedom. I shall say almost nothing more about this. There is little point either in defending the indefensible or in attacking a practice that

almost every reader here will already condemn. (Belshaw 2016, p. 11-12)

Such arguments therefore also entail a rejection of the acceptability of much current meat production and consumption.

Other defences are even more circumscribed. For example, Donald Bruckner (2016) argues that not only can meat eating be permissible, but that a consequentialist vegetarian has reason to eat fresh road-kill: because the animal is already dead, less harm will come from eating it than from eating plant-based foods that necessarily have environmental costs associated with production and transportation. Similarly, Singer and Mason (2007, p. 238-244) consider freeganism (salvaging food thrown away by supermarkets, usually from dumpsters) an ethical option, even where it includes meat, as it reduces waste and does not contribute to demand for animal products. Such approaches are extended by Bob Fischer (2018), who argues for the permissibility of eating road-kill, bivalves, and insects. While these arguments may show ethical ways of eating meat, they do not defend meat eating as it is commonly practised. They merely provide a limited way of eating meat that does not conflict with the first two beliefs of the meat trilemma.

More vigorous defences of meat eating face similar problems. For example, Roger Scruton (1998) argues against the concept of animal rights. Given that he establishes, to his satisfaction, the absurdity of the animal rights position, he argues that it is permissible to use animals for our purposes, including consumption that causes premature death. Nevertheless, he denies the acceptability of intensive animal agriculture, writing that:

Someone who was indifferent to the sight of pigs confined in batteries, who did not feel some instinctive need to pull down these walls and barriers and let in light and air, would have lost sight of what it is to be a living animal [... A] true morality of animal welfare ought to begin from the premise that this way of treating animals is *wrong*, even if legally permissible. (Scruton 1998, p. 102, emphasis added)

Thus again, even if Scruton succeeds, his argument does not establish the moral permissibility of eating intensively farmed meat.

The most sophisticated defence of meat eating I have come across is by Loren Lomasky. In a nuanced application of the *nice* justification discussed in Section 2.2.1.4, Lomasky (2013) unapologetically defends meat eating, arguing that the exquisite dining pleasures afforded by meat are an important (though neither necessary nor sufficient) factor in a life well lived. His point is not just that eating meat gives pleasure, but that the specific form of the pleasure can be morally

important, akin to the life-enriching pleasure of seeing a great and famous artwork in person or attending an excellent concert (Lomasky 2013, p. 193). He argues that because of the key role meat plays in life-enriching meals, eating meat, including that of intensively raised animals, can be morally permissible. According to Lomasky, meat's permissibility requires that the costs of animal suffering for meat production be balanced against the combined benefits of enjoyment to meat eaters and the pleasures that these animals would not have had without the life granted to them by the practice of meat eating. He claims that the practice of meat eating is beneficial to both humans and farmed animals, and thus animal suffering and death can be justified by the role meat plays in the good life. While Lomasky avoids interrogating what farming practices his standard would rule out, he thinks it can justify most if not all current farming practices.

However, even if Lomasky's moral calculus is correct, it does not justify anywhere near the levels of meat consumption found in most developed nations. Lomasky defends meat eating not just because it brings pleasure, or because (as he claims) meals containing meat provide more pleasure than vegetarian meals, but because meals containing meat allow the opportunity for unique and exquisite dining pleasures and valued aesthetic experiences. Yet many meat-containing meals, such as the cheap takeaway burger or the weeknight stir-fry, are expected to be ordinary or even mediocre. Lomasky (2013, p. 181) argues that such forgettable meals provide a necessary background against which exquisite meals can shine. Yet if it is the exquisite meat-containing meals that justify meat eating, then Lomasky's argument does nothing to justify the eating of meat-containing meals which can reasonably be expected to fall far short of that standard. The baseline against which exquisite meat-containing meals can be judged can just as easily be set by vegetarian meals as by meals containing meat. Thus while Lomasky's argument can, if it succeeds, justify eating exceptional meals containing meat (or ones that can be reasonably expected to have a good chance of being exceptional), it cannot justify the practice of everyday meat eating.

The majority of philosophical arguments for the permissibility of meat eating do not attempt to defend common agricultural practices. Instead, they explicitly rule out a defence of such practices and instead defend a much more circumscribed subset of meat. Lomasky's argument, while defending a greater range of practices, fails to defend ordinary everyday meals containing meat. Thus even if those arguments succeed, the scope of the 'good meat' they defend is far too narrow to enable good-conscience meat eating given current practices.

2.3.2 *Arguments that achieve too much*

Few philosophers defend contemporary animal agriculture – Timothy Hsiao’s (2015, 2017) work is the only direct defence of factory farming I have been able to find. Of the philosophical arguments which either directly or indirectly support contemporary animal agriculture, all require accepting premises unpalatable to most meat eaters. Specifically, they require a rejection of the first belief of the meat trilemma: they all accord *no* moral status to animals. Thus these arguments cannot be used with integrity by people caught in the meat trilemma unless those people are willing to revise their beliefs about the moral status of animals.

The most historically prominent position from which to defend contemporary meat eating is Descartes’. Descartes (1637, p. 139-141) argued that unlike people, animals are mere automata who are not sentient and so do not feel pain – they merely give the appearance of suffering. Thus acts which appear to be cruel or callous are morally permissible; thinking those acts cruel or callous is mere sentimentality. As Tom Regan (2004, §1.7-1.8) argues, such a position is extremely implausible given current scientific knowledge about the common evolutionary origins of humans and non-human animals – given our shared neural architecture, it makes no sense to think that the capacity to experience pain evolved only in humans. More deeply, however, irrespective of the possible merits of this argument, it has implications which most meat eaters would be unwilling to accept. Endorsing this argument requires accepting the position that hurting animals of any species is *never* wrong for what it does to the animal; it is a rejection of belief 1 of the meat trilemma. It requires the meat eater who uses this defence to hold that someone who is cruel to animals purely for enjoyment does no wrong (except, perhaps, to their own character), and this is not a bullet that many meat eaters will be prepared to bite.

The argument from contractarianism, such as that articulated by Peter Carruthers (1992, p. 194-196), faces a similar problem. Carruthers does not deny that animals can feel pain, but instead argues that because animals are not rational agents, they are not part of the social contract.¹⁰ Thus according to Carruthers, we have only indirect duties toward animals – our only duties to them stem from a respect for the sympathies of animal lovers, or from a concern for the effects good or bad treatment of animals may have on our character (Carruthers 1992, p. 194). As such, “there is no basis for extending moral protection to animals beyond that which is already provided. In particular, there are no good moral grounds for forbidding hunting, factory farming, or laboratory testing on animals” (Carruthers 1992, p. 194). On this approach, affection or sympathy can ground care toward an animal,

¹⁰ Contractarianism need not be hostile to non-rational agents, see, for example, VanDeVeer (1979).

but this care is not owed to the animal. Hence Carruthers reaches the same position as Descartes, albeit for different reasons: animals, in and of themselves, fall entirely outside the scope of morality. Endorsing his argument therefore again requires the meat eater to bite the same unpalatable bullet: hurting animals is never wrong *because of what it does to the animal*, only because of effects it may have on people.

Timothy Hsiao (2015, 2017) presents the only direct and explicit defence of intensive animal agriculture I have been able to locate, and his argument faces similar problems. Hsiao accepts that industrial farming practices inflict significant pain on and cause harm to animals. However, he argues that the harm involved in killing a non-human animal is analogous to the harm involved in killing a rose bush – neither is morally significant (Hsiao 2017, p. 41). Unlike Descartes, he grants that animals have a mental life, but he denies that simply having a mental life is sufficient for moral status. Instead, he argues that a capacity to reason, rather than merely having a welfare, is what grounds moral status.¹¹ Causing pain or harm is therefore only wrongful – only morally significant – when it occurs in beings “capable of rational agency” (Hsiao 2015, p. 279-280) and thus “animal suffering is not morally salient [because] animals lack the required features necessary for membership in the moral community” (Hsiao 2017, p. 38). Like Carruthers’ argument, Hsiao’s requires establishing that *all* non-human animals lack moral status and hence that there are no limits on what can permissibly be done to animals, except where this affects the sensibilities of others or causes harm to a person’s character. While this approach may, if all premises are accepted, justify meat eating as currently practised, the cost is one which most meat eaters, particularly those who identify as animal lovers, will be unwilling to pay.

The literature in defence of meat eating does not, for the most part defend meat eating as it is currently practised. Most arguments, even if they succeed, do not succeed in defending most of the meat that is eaten – in particular most chicken and pork. Those defences that do not exclude the meat of ‘unhappy animals’ are few, and require holding that animals have no intrinsic moral worth. Rejecting the belief

¹¹ Since many humans (including infants and people with severe cognitive disabilities) lack the capacity to reason, he suggests that the capacity to reason which matters is not actualized capacity to reason but rather a root capacity for reason. This root capacity is claimed to exist in all human beings but in no non-human animals. However, Hsiao (2015, p. 288) is not committed to this, acknowledging that it may instead be that people who cannot reason also lack moral status. It is also plausible that someone following a similar line of argument could draw the line of having sufficient reasoning capacity for moral status elsewhere, including for example the great apes or dolphins while excluding less ‘intelligent’ animals. However, it is highly implausible that such an approach could draw boundaries where the average meat eater wants them. In particular, it would be difficult to draw the boundary to include dogs and cats within the moral sphere while excluding pigs and cows.

that animals have moral value is something that few people currently caught in the meat trilemma would be willing to do, and so these arguments cannot do the needed justificatory work.

2.4 MAINTAINING GOOD-CONSCIENCE MEAT EATING

Thus far in this chapter, I have established that for most meat eaters, their belief in the moral acceptability of meat eating is inconsistent with their other beliefs. Nevertheless, most meat eaters eat meat in good conscience and continue to feel confident that meat eating is morally acceptable even when confronted with the shortcomings of the available justifications. This section presents a sketch, to be extended in the following chapters, of why this is. I argue that justifications like the ones described above, along with *strategic ignorance*, play an important role in managing *cognitive dissonance* – they make meat eating seem excusable even if they fail as justifications. Furthermore, I argue that they succeed in this function because meat eating is so widespread. Meat’s widespread acceptance also makes it easy to treat the animal suffering involved in meat production as being merely bad rather than morally wrong. This distinction between bad and wrong also helps with cognitive dissonance management.

Eating meat is, for most people, enjoyable. However, this enjoyment comes with psychological costs. Contemplation of animal suffering is unpleasant, as is the idea that in eating meat one might be morally responsible for at least some of this suffering. The disconnect between the enjoyment of meat and the discomfort with animal suffering can give rise, in many meat eaters, to the unpleasant feeling of cognitive dissonance. Because cognitive dissonance is unpleasant, people are motivated to reduce it. According to Matthew Rabin (1994, p. 178), cognitive dissonance can be reduced in one of two ways: modifying one’s behaviour (in this case, by limiting meat consumption to ‘ethical’ choices, or eliminating it altogether), or modifying one’s beliefs. Justifications like those described above are an attempt to do the latter, and they work.

Such justifications can help alleviate cognitive dissonance. Meat eaters who endorse the 4Ns feel less guilt about their meat eating than people who avoid thinking about the issue (Piazza *et al.* 2015, p. 123). Furthermore, guilt about dietary choices negatively correlates with strength of endorsement of these justifications (Piazza *et al.* 2015, p. 123). Thus irrespective of the philosophical rigour of these arguments, they play a role in making meat seem morally acceptable – they alleviate cognitive dissonance. However, because of their lack of rigour, there is a psychic cost to maintaining those beliefs:

In general, there is likely to be a natural, intellectually honest set of beliefs about the morality of an activity. Developing beliefs that differ from this level is costly because

it may intrinsically conflict with other parts of a person's belief system, and reintegrating it can involve laborious intellectual activity. (Rabin 1994, p. 180).

Justifications like those discussed above can therefore help manage cognitive dissonance. However, doing so imposes the cost of seeking out these justifications and buttressing them as necessary when they are challenged. This cost is inversely proportional to how widespread acceptance of these justifications is. The more widely accepted a position or behaviour, the easier it is to convince oneself that it is morally acceptable. Conversely, if nobody else thinks an activity is ethical, then it is hard to convince oneself that it is (Rabin 1994, p. 179). Justifications can therefore help manage the cognitive dissonance arising from meat consumption, but the prevalence of meat eating is important to enabling this function.

This may be because, as Mercier and Sperber (2011) argue, reasoning primarily evolved to serve an argumentative role. Reasoning is therefore primarily used to bolster an existing position – people tend to think of reasons to support what they already believe, and rely on others to highlight weaknesses in their position rather than thinking of counter-arguments themselves (Mercier and Sperber 2011, p. 65). Thus as long as a position is not challenged, weak arguments in favour of that position can survive without being scrutinized. Conversely, where a position is likely to be disputed, one will be faced with counter-arguments which need to be addressed.

Hence an important mechanism for maintaining dissonant beliefs is avoiding unwanted information – both information that would undermine one's justifications and information that would make salient the moral implications of what one is doing. In cases of *strategic ignorance*, people avoid information that might cause them to have to change their behaviour. Strategic ignorance is therefore a way of keeping the costs of maintaining intellectually dishonest beliefs low. For example, Onwezen and Weele (2016, p. 97) demonstrated experimentally that people are less inclined to behave selfishly if informed how their choices will affect others. However, given the choice, a significant proportion of people will decline to be informed (i.e. adopt strategic ignorance) and continue to “engage in personally maximizing strategies” (Onwezen and Weele 2016, p. 97). Strategic ignorance thus enables people to make selfish choices while maintaining plausible deniability (both to themselves and to others) about having done so.

A clear example of strategic ignorance can be observed in the results of a study by Dana, Weber, and Kuang (2007). They found that in a dictator game variant where players had the choice between keeping \$5 for themselves while giving \$5 to the other participant, or keeping \$6 for themselves while giving \$1 to the other, 74% chose the fair division. However, self-maximizing choices became far more com-

mon when the possibility of strategic ignorance was introduced. In the variant allowing strategic ignorance, players again had the choice between receiving \$5 or \$6, with either \$1 or \$5 going to the receiver. However, whether the pairing was \$5-\$1, \$6-\$5 or \$5-\$5, \$6-\$1 was obscured unless the dictator chose to gain this information. In this variant, 46% of participants chose to remain strategically ignorant and receive \$6, while only 47% chose to learn how much the receiver would gain and then chose the fair option. Without knowledge of the recipient's payoffs, senders had a 50% chance of choosing the option most beneficial to the recipient while also choosing the option most beneficial to themselves; with knowledge of the recipient's payoffs, senders risked being forced, by feelings of empathy or guilt, to choose the option less beneficial to themselves. Hiding behind ignorance enabled players to maximize their payoffs without seeing themselves as bad (Dana 2005, p. 211). Thus many players chose to remain ignorant of how their choice would affect the other player.

These two studies reveal that foregoing even freely obtainable information about harms to others (strategic ignorance) can enable self-serving misconduct. Instead of directly violating moral rules (which suggests malicious intent), one can instead circumvent them (Dana 2005, p. 206). Strategic ignorance thus allows a sufficient level of plausible deniability – one can make self-interested decisions in good conscience even where information that one should act otherwise is readily available.

Strategic ignorance can help avoid making both the suffering of animals and the inadequacy of justifications salient. Self-maximising behaviour (in the form of eating meat) can be enabled by strategically ignoring information about the conditions under which farmed animals live, for example by turning away from undercover videos or refusing to read vegetarian literature. Experimental research by Onwezen and Weele (2016, p. 99) confirms this – they found that many consumers admit to ignoring information about the experiences of fast-growing chickens and avoiding thinking about these possibilities when buying chicken meat. Furthermore, this strategy works. Onwezen and Weele (2016, p. 100) found that “a higher score on willingness to ignore is associated with experiencing fewer negative emotions and less perceived responsibility.” Such strategically ignorant meat eaters:

experience low amounts of negative emotions. These consumers score high on the willingness to ignore and appear to strategically ignore the issue. They feel responsible themselves, but they also feel that others are responsible. The results imply that these consumers feel responsible but strategically ignore the issue. As a result, they experience low levels of cognitive dissonance. (Onwezen and Weele 2016, p. 102)

Thus strategically ignorant meat eaters can avoid contemplating whether or not meat eating should be treated as a moral issue. In avoiding the issue, they avoid the need to revise either their beliefs or their dietary habits.

Strategic ignorance can also help maintain justifications. Belief in the necessity of meat eating might be maintained by avoiding information about the nutritional possibilities of a vegetarian diet, especially as seeking out such information requires some effort. Likewise, the belief that normalcy, naturalness, and niceness justify meat eating can be maintained by avoiding thinking about the problems with those justifications. Strategic ignorance is therefore an effective strategy for keeping meat-related cognitive dissonance under control.

However, strategic ignorance is only possible where ignorance is plausibly possible. The more widely available information on a given topic is, the higher the effort required to remain ignorant (Rabin 1994, p. 179). Thus while strategic ignorance about meat is a tactic adopted by individuals, it relies on not being extensively challenged to succeed. A knowledge threshold exists – once this threshold is exceeded, not knowing becomes seen as culpable and blameworthy rather than as an excuse. For most meat eaters, those to whom my conditions apply, good-conscience meat eating therefore requires that meat eating be a widespread practice. The justifications discussed in [Section 2.2](#) could not alleviate cognitive dissonance if they were not widely accepted, and strategic ignorance would be much harder to maintain if most people were well informed.

Another important factor is perceived excusability. Many people agree that it would be morally good to stop eating meat while denying that they have a moral obligation to do so – they treat vegetarianism as being supererogatory. We all fail to act as well as we should in some areas – we purchase cheap and disposable products that are environmentally irresponsible, use too much energy, fail to donate to charities, and eat meat. Yet always choosing the morally good option is extremely demanding, and few people aspire to live up to the utilitarian ideal of always choosing the option that maximises the good. Given that we see ‘good’ people who fail to act on these issues, it seems consistent that one can be a good person while also failing to always do the right thing, and that one can even do morally bad things without thereby being a bad person.

Intuitively, some moral failings are excusable. Given that always choosing the morally better option seems too onerous and failing can be excusable, it is not clear that meat should be singled out as demanding action. Because meat eating is widespread, it intuitively seems to belong in the category of excusable wrongs. It may be hard to justify being a good person while going on murderous sprees – such things simply are not done. However, it is much easier to justify that one is a generally good person while manifesting a common

moral failing such as eating meat. This explanation again points to the importance of what other people do. If meat eating were uncommon, it would be difficult to justify one's meat eating in the same way – it is difficult to see oneself as a good person who occasionally does bad things when the occasional bad one commits is one deeply condemned by one's society. I return to this idea in [Chapter 4](#) when I explore the role other people's behaviour has in shaping how meat eating is understood.

2.5 CONCLUSION

When ethical dilemmas are discussed, whether by philosophers or by laypeople, any number of factors tend to be weighted for and against different options. However, the factors most commonly taken to support meat eating are not seen as having much if any weight in 'real' moral issues.

Of the blocking justifications discussed, only strong necessity can make acceptable something that would otherwise be morally wrong; weak necessity, naturalness, normalcy and niceness merely give one reasons in favour of something morally neutral. Yet most meat eaters could gain adequate nutrition from a meat-free diet; for them, the strong necessity claim is false, and patently so given the existence of healthy long-term vegetarians. Justifications from weak necessity, naturalness, normalcy and niceness also fail to resolve the meat trilemma. The effort involved in avoiding meat is insufficient for weak necessity to do the necessary work. We do not take the naturalness of something to have independent moral value – only natural things that we also approve of are defended in that manner. Likewise, we do not think that historical practices such as slavery or the practices of other cultures such as female genital cutting were or are justified by their normalcy. Nor do we take potential pleasure to justify harming others who have moral status. The use of such justifications requires that animals do not have moral status, yet the moral status of animals, including animals of species commonly killed for meat, is affirmed by those to whom the meat trilemma applies. The no-effect justification fares similarly – many people who utilize it have strong feelings about the importance of integrity or complicity with wrongdoing, at least in other domains. Hence this justification can only rescue the strict consequentialist. Yet most people caught in the meat trilemma follow a folk morality that involves many deontological considerations, and so their invocation of the no-effect defence is ad hoc.

The more sophisticated philosophical arguments in favour of meat eating do not fare any better. Most have an extremely limited scope: they aim to justify eating meat where it would otherwise 'go to waste', or only of animals who have lived good lives. Such arguments justify meat eating for people prepared to be extremely selective about the

meat they eat, but they cannot help the meat eater who wishes, in good conscience, to continue eating the purchased meat of intensively farmed animals – to continue eating meat that is cheap and easy to source. Others try to justify meat eating on its unique aesthetic value, but this fails to defend ordinary meat consumption. Philosophical arguments that can be used to defend the everyday consumption of intensively farmed meat require that the endorser of those arguments give up the belief that animals have moral status. While this might be endorsed by some meat eaters, most will find this repugnant and shy away. Such people cannot justify their meat eating through reference to these arguments.

Most people's meat eating therefore cannot be justified without the revision of important beliefs about the general permissibility of causing animals to suffer, or about how one should respond to moral demands. Hence the available justifications cannot explain away the *prima facie* moral badness of meat eating. Yet the justifications I have discussed are resistant to being dispelled. This is because they play an important role in managing the cognitive dissonance that arises in meat eaters who also care about animals. To manage cognitive dissonance, people are motivated to maintain these beliefs even in the face of contrary evidence. Along with strategic ignorance (deliberately avoiding inconvenient information) these justifications enable good-conscience meat eating. However, the success of these strategies depends on their being widely accepted – otherwise the effort to maintain beliefs about the justifiability of meat eating would be much greater.

Meat eating provides an interesting moral case study – it is unlike prototypical moral cases like theft and murder which are tempting only in extreme circumstances. Instead, eating meat is something which many people either have some doubts about or explicitly think immoral, but nevertheless strongly desire to do. It tends not to be justified as a regrettable but necessary evil, but rather to be embraced enthusiastically. My analysis of the common justifications for eating meat highlights that most meat eaters begin from the position that eating meat is not a moral issue; the justifications they employ would not be accepted by those same people if applied to something thought to be immoral. Rather, their perceived justificatory acceptability presupposes that meat eating is not morally wrong. Because of this, as I will argue in [Chapter 3](#), reasoning with someone is unlikely to persuade them to change their behaviour. At most, it will turn them into a *conflicted carnivore*, someone who is persuaded that meat eating is immoral but does it anyway. The widespread acceptance of meat eating makes it seem excusable even if not strictly justifiable – thus a meat eater can think they do something *bad* in eating meat without thereby doing something *wrong*. In other words, badness does not straightforwardly reduce to wrongness, and the meat eater implicitly

exploits this argumentative lacuna to eat meat in good conscience. In [Chapter 4](#), I therefore explore how social norms shape what is likely to be perceived as merely bad, and what is likely to be perceived as wrong.

THE PUZZLE OF THE CONFLICTED CARNIVORE

3.1 MEAT IS BAD, PASS ME THE CHICKEN

When confronted with an unambiguously moral issue, people tend to be highly motivated to comply with what morality requires. Under ordinary circumstances, for example, it is easy not to steal even if we badly want a thing that we cannot otherwise obtain. Similarly, it is easy for most people to remain motivated not to murder another person, even if they are making one's life extremely difficult. Indeed, the very idea of being forced by circumstance to violate one's morals by stealing or killing is upsetting. Abstention from meat, however, tends to be hard to motivate, even for most people who claim to think meat production and consumption a serious moral issue. In this chapter, I seek to understand the difference between strongly motivating moral issues, and issues where, as with meat eating, it is easy to turn away from the moral implications of what one does. To this end, I focus on *conflicted carnivores*: people who admit that the justifications described in [Chapter 2](#) fail yet continue to eat meat, mostly in good conscience. Their explicit acknowledgement of the immorality of their diet generates an interesting puzzle: how is it that they can say that animal suffering for meat production is a moral issue, yet find it easy to turn away? I argue that a strongly motivating judgement requires both a rule endorsed by a person (and usually by their society) and an affective response to violations of that rule. Given that we do not have rules against eating meat, the first component is not present and so affective responses, even if experienced, tend not to be treated as carrying morally important information.

Although it can be readily observed that people do not consistently do what they profess to think is the right thing, this should nevertheless be a puzzling rather than a mundane fact about human behaviour. As Michael Smith (1994, p. 6) highlights, it simply does not make sense to say "Yes, you have persuaded me that I should ϕ , and I have no extenuating circumstances, but nevertheless, I won't do it." Yet this is precisely what the conflicted carnivore does. Such behaviour seems to exhibit not just a lack of moral fibre, but also a failure of rationality: it seems that such a person has failed to understand what making a moral judgement entails.

It might therefore seem that the conflicted carnivore's behaviour can be explained semantically, but such an answer is philosophically uninteresting. Many people (moral judgement internalists) deny that a moral judgement is being made at all in cases like that of the con-

flicted carnivore. For simplicity, I follow ordinary language usage in calling the conflicted carnivore's judgement a moral one. However, my argument does not rely on this terminological choice, and readers who object can replace "moral judgement" with "moral-like judgement". My interest is in *why* the conflicted carnivore's judgement against meat is not motivating, not in whether it should be called a moral judgement proper. I argue that irrespective of the meta-ethical position one takes regarding the connection between moral judgement and motivation, the fact that the conflicted carnivore is not motivated to comply with their judgement about meat, but is motivated to comply with other 'moral' judgements, is a discrepancy that needs explaining. No meta-ethical position on what constitutes a moral judgement can explain *why* it is that the conflicted carnivore is not motivated to stop eating meat but is motivated to comply with at least some moral judgements.

I propose that a first step towards understanding the conflicted carnivore's behaviour can be made by applying Paul Rozin's idea of *moralization*. When a judgement is moralized, it becomes internalized such that acting against it feels deeply wrong. In other words, moralization is the process by which a judgement is transformed from one which would be recognized as a moral judgement by externalists (but not internalists) to one that would be recognized by both externalists and internalists. A conflicted carnivore, by definition, has not moralized their judgement against meat. Conversely, a struggling vegetarian has partially moralized their vegetarian position, while the vegetarian to whom the very idea of consuming animal flesh is abhorrent has more completely moralized their judgement against meat. However, using moralization to provide an answer again generates a question: what causes moralization, and why is it only sometimes triggered in the case of meat?

To answer this, I canvass a number of psychologically informed theories of moral judgement. I first argue that people cannot consistently reason their way into making and acting on moral judgements based on abstract intellectual arguments alone. I then turn to sentimentalist theories, which are better able to account for cases where a meat eater fails to be affectively moved by meat. Nevertheless, I argue they are unable to offer a sufficient explanation for why deeply distressing affective stimuli so often fail to provoke a moralized response. I conclude by arguing that the under-determinacy of affective responses can be explained by Shaun Nichols's (2004) *Sentimental Rules* framework: something tends to be treated as properly moral (i.e. moralized) where there is both a negative affective response to that thing *and* an endorsed rule against the activity that generates the affective response. Although this rule can be a strongly internalized personal rule, in most cases, effective rules will be socially endorsed ones. Hence, the behaviour of others will have a strong influence on

the likelihood of an intellectual judgement about a moral issue becoming moralized, as will be more thoroughly explored in [Chapter 4](#).

3.2 THE CONFLICTED CARNIVORE

Conflicted carnivores are people who explicitly endorse the idea that they morally ought not eat meat, but nevertheless enjoy eating it and experience little guilt.¹ In my study with John Thrasher, 55% of survey respondents agreed that meat eating is morally unacceptable, but nevertheless reported liking and eating meat (Thrasher and Mayne 2018). While such people rationally reject pro-meat arguments like those discussed in [Chapter 2](#), their behaviour stands in contrast: both behaviourally and emotionally, meat eating seems to be ‘not a big deal’. They may feel an occasional pang of guilt about their meat eating, but for the most part, the moral conflict is not serious.

The tension between how conflicted carnivores think they should feel and act and how they actually feel and act is nicely captured in the following reflection by social psychologist Jonathan Haidt, a continuing conflicted carnivore:

Singer’s approach to the ethics of killing animals changed forever my thinking about my food choices [...] Singer’s clear and compelling arguments [in *Animal Liberation*] convinced me on the spot, and since that day I have been morally opposed to all forms of factory farming. Morally opposed, but not behaviorally opposed. I love the taste of meat, and the only thing that changed [...] after reading Singer is that I thought about my hypocrisy each time I ordered a hamburger. (Haidt 2006, p. 165)

A similar tension was experienced by novelist Jonathan Safran Foer, whose lasting vegetarianism was reached only after years of alternating between vegetarianism, guilty meat eating, and enthusiastic meat eating:

What our babysitter said made sense to me, not only because it seemed true, but because it was the extension to food of everything my parents had taught me. We don’t hurt family members. We don’t hurt friends or strangers [...] I had to change my life. Until I didn’t. My vegetarianism, so bombastic and unyielding in the beginning, lasted a few years, stuttered, and quietly died. I never thought of a response to our babysitter’s code, but found ways

¹ Strictly speaking, such people are omnivores, not carnivores. All humans, even those who adhere to a strict vegan diet, are technically omnivores, as ‘omnivore’ is a biological designation regarding what one is able to digest, not what one chooses to eat (see Joy (2010, p. 30) for further discussion). However, the term ‘carnivore’ readily conveys the idea of a meat eater and allows for a catchy category name.

to smudge, diminish, and forget it. Generally speaking, I didn't cause hurt. Generally speaking, I strove to do the right thing. Generally speaking, my conscience was clear enough. Pass the chicken, I'm *starving*. (Foer 2009, p. 7, emphasis original)

Prima facie, it seems that conflicted carnivores like Haidt and the young Foer are making a moral judgement against meat eating, but for some reason, this moral judgement is failing to be manifested, or to be manifested in a sustained way, in their actions.² Thus conflicted carnivores pose an interesting puzzle – they endorse through their behaviour something they claim to disavow. They are people who seem to have been persuaded, yet simultaneously remain un-persuaded, by the arguments for vegetarianism.³

3.3 A PUZZLE ON ANY VIEW OF MORALITY

Conflicted carnivores seem to have made a moral judgement against meat eating: they admit that eating meat is morally wrong. This judgement, however, fails to motivate them to stop eating meat. This motivational deficit needs to be explained, and it needs explaining irrespective of how one views the connection between judgement and motivation.

There are four ways to understand the connection between a person's moral judgement and their motivation to act on it:

1. **Strong internalism:** If one makes a moral judgement then one is necessarily and overridingly motivated to act in accordance with that moral judgement; being motivated is a definitional component of making a moral judgement (e.g. Mackie's (1977) reading of Plato).
2. **Weak internalism:** If one makes a moral judgement and is practically rational (i.e. not suffering from weakness of will or a related phenomenon), then one is necessarily motivated to act in accordance with that moral judgement (e.g. Smith (1994)).
3. **Generic connection (externalist):** When people make moral judgements they tend, perhaps due to deep facts about human nature, to be motivated to act in accordance with that judgement, but this motivation is not conceptually necessary (e.g. Hume (1740)).

² While some conflicted carnivores reduce meat consumption or avoid the most egregious of animal products such as veal or *foie gras*, such reduction is insufficient to resolve the puzzle. Such people nevertheless eat meat, and when they do, they enjoy doing something they consider to be wrong.

³ As discussed in Chapter 1, this form of inconsistency extends well beyond meat eating, but I limit my focus to this case study.

4. **Contingent connection (externalist):** The connection between moral judgements and motivation is merely contingent on one or more other factors; people will sometimes be motivated by moral judgements, and sometimes they won't (e.g. Brink (1989)).

Each of these positions might seem to have implications for how the conflicted carnivore's behaviour is to be understood. However, a closer inspection reveals that while each gives an account of how much, if at all, the motivational gap should matter for how the conflicted carnivore's judgement is understood, none of them explains the cause of this gap.

3.3.1 *Internalist positions*

If there is a necessary connection between making a moral judgement and being motivated to act on it, then it might seem that the conflicted carnivore's motivational gap can be readily explained as stemming from insincerity of judgement or weakness of will. However, such explanations fail to address the deeper and more interesting questions of why some judgements tend to be sincere while others are not and why judgements about meat generally fall in the latter category.

Both strong and weak internalism hold that it is conceptually necessary for a moral judgement to be motivating. Strong internalism is the conceptual claim that necessarily, if one sincerely judges that one ought to ϕ , then one will be overridingly motivated to ϕ without needing any accompanying desires (Rosati 2014, §3.2). Hence it is not possible to judge that one should ϕ without thereby being overridingly motivated to ϕ . Under this conception of morality, a judgement that is not complied with is not sincere and is therefore not a moral judgement. Hence the conflicted carnivore's judgement against meat eating simply is not a proper moral judgement.

Conversely, weak internalism holds that when one makes a moral judgement, one necessarily gains some motive to act in accordance with that moral judgement (Rosati 2014, §3.2). Under weak internalism, a moral judgement must necessarily be at least somewhat motivating, though this motivation need not be sufficient to generate compliance. Hence, if one judges that it is morally wrong to ϕ , this necessarily gives one at least some motivation not to ϕ . Similarly, if one judges that it is morally right to ϕ , then one will necessarily be at least somewhat motivated to ϕ . This understanding of moral judgements allows that the conflicted carnivore's judgement may be sincere but insufficiently motivating.

The strong internalist understanding of moral judgement seemingly provides a straightforward explanation for the conflicted carnivore's judgement-behaviour discrepancy: the conflicted carnivore simply has not made a genuine and sincere moral judgement. This

explanation might also sometimes apply on the weak internalist account. However, this analysis of the situation can be accepted without it providing an interesting answer to the puzzle. An explanation is needed at a lower level – the question becomes “What is it about sincere moral judgements that distinguishes them from insincere ones, and why are the conflicted carnivore’s judgements of the insincere variety?” Merely categorizing the judgement as insincere provides no explanation: it does not say anything about why the judgement is insincere, nor what it would take to convert it to a sincere moral judgement.

On the weak internalist view, there is also another possibility: it could be that a sincere moral judgement is being made, but that the conflicted carnivore is “weak willed”. Weakness of will is an intuitively plausible explanation for a variety of moral failings, yet it fails to provide a helpful explanation for the conflicted carnivore’s apparent motivational deficit. People commonly exhibit weak will in certain moral (and non-moral) domains while not struggling in others. A person might, for example, cheerfully eat meat all their life despite thinking this to be wrong, while never stealing or even being seriously tempted to. This is particularly interesting in cases where a conflicted carnivore might, on reflection, admit they think eating meat to be morally worse than petty theft, while being strongly motivated to avoid the theft, but not the meat.⁴ Thus an explanation from weakness of will again raises a deeper and more interesting question: “Why do conflicted carnivores (and some avowed ethical eaters) suffer from weakness of will regarding meat when they have little or no difficulty complying with other moral demands?” Weakness of will does not explain the motivational deficit. Rather, it highlights some of the underlying factors a fuller explanation will need to account for. Thus the puzzle of the conflicted carnivore remains a puzzle on internalist understandings of morality.

3.3.2 *Externalist positions*

Externalists deny “that it is a conceptual platitude that necessarily moral judgments motivate” (Rosati 2014, §3.2). Rather, a moral judgement is motivating only if it is accompanied by a desire that causes one to want to comply. On such a view it is unsurprising that conflicted carnivores are not motivated by their judgement against meat – clearly, the necessary desire is missing. Nevertheless, there remains a question of what it is that makes some moral judgements motivating, and why this factor is absent for the conflicted carnivore.

Explaining the conflicted carnivore’s motivational gap as stemming from the lack of a motivating desire is again too shallow. The first externalist way of understanding the connection between judgement

⁴ Thanks to Shaun Nichols for helping me draw this point out more clearly.

and motivation is as a generic connection: the prototypical moral judgement is motivating, but any given moral judgement need not be motivating. The other possibility is that the connection between judgement and motivation is purely contingent: some moral judgements are accompanied by motivation to comply, some are not. On both these conceptions, one can say that moral judgements need not motivate and that the conflicted carnivore's judgement just is of the non-motivating kind. However, this is not an interesting explanation as it does not indicate what factor or factors transform mere moral judgements into motivating ones, nor why these factors are absent for the conflicted carnivore but present for vegetarians and vegans. An interesting explanation should also say something about what would cause the conflicted carnivore's mere moral judgement to be transformed into a motivating one. Thus the conflicted carnivore's judgement-motivation gap also remains a puzzle on externalist conceptions of morality.

3.3.3 *The common threads*

The above discussion reveals that despite initial appearances, the internalism/externalism debate is ultimately orthogonal to answering the puzzle of the conflicted carnivore. None of the discussed understandings of the connection (or lack thereof) between judgement and motivation can provide a sufficient explanation of the conflicted carnivore's motivational gap. If one rejects that the conflicted carnivore has made a genuine moral judgement (i.e. if one holds that the conflicted carnivore does not *really* think that eating meat is wrong), then either something must be blocking this judgement, or some key ingredient must be missing, and an interesting answer to the puzzle must say what this is. Similarly, if one thinks that weak internalism is true and that the conflicted carnivore is making a genuine and somewhat motivating moral judgement, then why is the motivation generated by the conflicted carnivore's judgement so weak? Alternatively, if one thinks that there is no necessary connection between making a moral judgement and being motivated to comply, and that the conflicted carnivore's judgement is simply missing the extra factor necessary to generate motivation, what *is* this missing factor and what would it take for this factor to be present? Irrespective of how one thinks the connection between moral judgement and moral motivation is best explained, the conflicted carnivore's motivational gap still needs an explanation. The remainder of this chapter will provide the beginnings of such an explanation, one which can be accepted irrespective of the position one takes on the meta-ethical possibilities discussed above.

3.4 COMPONENTS OF MORALIZED JUDGEMENT

To understand, at a deeper level, why someone would eat meat despite thinking it wrong, we must understand the psychology of moral judgement and how it connects to moral motivation. Not all moral issues are experientially alike. Some moral behaviour is automatic and, for the most part, easy to comply with. Other moral demands are, subjectively, demanding. For example, refraining from stealing comes with costs, especially where the risk of being caught is low. However, for most people, most of the time, stealing is simply unthinkable. Similarly, most people coming across a lost child would feel compelled to stop and help. In these cases, behaving morally requires little or no willpower. Rather, effort would be required to ignore one's inner moral promptings and follow through on the immoral choice.

In other cases, meat eating being a prime example, the moral calculus for most people does not lead to judgement-compliant action. Although conflicted carnivores say they consider meat eating to be immoral, eating meat doesn't *feel* wrong to them. Thus their intellectual judgement against meat does little or nothing to influence their behaviour. While meat eating may, intellectually, seem worse than other wrongful behaviours such as petty theft, it is much harder to resist.

The thinkability of meat can be a problem even for people who have committed to a meat-free diet. A recent study by the Humane Research Council (2014) found that 84% of vegetarians and vegans abandon the diet (p. 4), while 16% of current vegetarians and vegans admit to having lapsed at some point (p. 9).⁵ Willpower can be required to stop eating meat and the will sometimes fails.⁶ For example, Marta Zaraska writes:

Sometimes, if no one can see me—and this is really difficult to admit—I nibble on a slice of sausage or a strip of bacon. It doesn't happen often—maybe once every six months or so. The taste usually disappoints me. I feel guilty over harming the poor cow, pig, or chicken and swear I'll never do it again. And then, sure enough, I do it again [...] There is something in it [...] that keeps luring me back. (Zaraska 2016, p. 5)

⁵ Specifically, 84% of people who have been vegetarian or vegan at some point in their lives no longer are. One third of former vegetarians maintained the change for three months or less (Humane Research Council 2014, p. 5). It is unlikely that such people were able to develop strong vegetarian habits; eating meat would have remained highly thinkable for them.

⁶ Though this study did not distinguish between those who are vegetarians for ethical reasons and those following a vegetarian diet primarily for health or weight-loss reasons, many ethical vegetarians struggle to consistently maintain a strict vegetarian diet.

Yet for other vegetarians and vegans, the very idea of consuming animal flesh is repellent; willpower has no role to play.

The difference between the meat-tempted and the meat-repelled vegetarian highlights that ‘moral’ judgements can, *prima facie*, arise in different ways. They can stem from an intuitive and affect-backed response which need not be rationally endorsed, or can be the product of an intellectual, rationally endorsed process. In the literature, the former are described as System I judgements, the latter as System II judgements (Stanovich and West 2000, p. 658).⁷ In everyday discourse, both judgements generated by System I and System II can be termed ‘moral judgements’. However, they track two different understandings of morality. There is a morality that is easy to comply with because it is less demanding and is accompanied by strong internal motivation (judgements generated by System I), and there is a morality that is difficult to live up to (the realm of System II judgements unsupported by System I).

Much of the moral philosophy literature is concerned with the debate about what we should consider ‘true moral judgement’. As Jeanette Kennett and Cordelia Fine (2009, p. 80) note, where inputs into a moral decision conflict, the ‘real’ moral judgement can be understood in different ways. It can be seen as the one most influential on action, the one that best meets our concept of what a moral judgement is, or the one which has normative authority. I want to sidestep this debate and leave open the conceptual question of what should be termed a *genuine* moral judgement. Instead, I focus on judgements about potentially moral issues. This includes both issues that are unambiguously immoral (for example, murder), and issues, such as meat eating, where this is less clear. I wish to explore questions surrounding how, and to what extent, motivation is generated (or not) by the different inputs into what are colloquially referred to as ‘moral judgements’. As such, I conform to colloquial usage, using the term *moral judgement* inclusively for both action-guiding and non-action-guiding judgements about what morality requires. However, as argued in Section 3.3, the ensuing discussion should be of interest even to someone who rejects the use of the term ‘moral judgement’ for non-motivating judgements.

3.4.1 Moralization

Paul Rozin’s 1999; 1997 idea of *moralization* is useful for capturing the difference between those moral judgements that are deeply motivating and those that are not. According to Rozin (1999, p. 218), “Moralization is the process through which preferences are converted into values”. A preference is seen as morally neutral (Rozin 1999, p.

⁷ As will be discussed below, a judgement that is initially the product of System II can over time also become a System I judgement.

220); it is something about which others can legitimately disagree, and something about which one can easily change one's mind. However, when a preference is transformed into a value, it becomes generalized. When moralization occurs, something one personally takes a con-attitude towards instead becomes seen as something which everyone should disapprove of.

Values are more motivating than preferences. Unlike preferences, values become an important part of one's self-conception; they are internalized and treated as stable and central to the self (Rozin, Markwith, and Stoess 1997, p. 67). A value is emotionally laden; the idea of violating it tends to be accompanied by disgust. While both preferences and values are linked to affective systems, violations of values tend to induce "strong moral emotions, such as anger, contempt, disgust, guilt, and shame" (Rozin, Markwith, and Stoess 1997, p. 67). A preference not to eat meat might be readily outweighed by other preferences. In contrast, a value of not eating meat will be much harder to override since the idea of violating a value is deeply unpleasant. A *moralized* judgement will always be motivating (though not necessarily overridingly so); a merely *moral* judgement, as I am using the term, need not be motivating. In other words, a moralized judgement is a judgement that a weak internalist would recognize as a moral judgement.

Moralization requires re-categorization. According to Rozin (1999, p. 219), moralization on an individual level occurs where new information induces someone to realize that an activity, previously seen as morally neutral, actually falls under the purview of an existing moral belief. This information can be cognitive (e.g. learning that gelatin is an animal product) or affective (e.g. seeing a fish gasping for air). The new information causes one to begin responding to the re-categorized activity in the same way as other activities covered by that moral belief. For example, a pescatarian might reject eating fish and become a vegetarian after the affective salience of the sight of fish struggling for air brings fish under that person's existing rule of 'don't eat dead animals' (Rozin 1999, p. 219). This "moral piggybacking" can be seen in Foer's quote from Section 3.2:

What our babysitter said [about killing and eating animals] made sense to me, not only because it seemed true, but because it was the extension to food of everything my parents had taught me. We don't hurt family members. We don't hurt friends or strangers. (Foer 2009, p. 7)

Foer's realisation that he should not eat animals piggybacked on his belief that causing harm is wrong. Animals, previously outside his moral sphere, were brought into it, albeit imperfectly. Moralization, however, was incomplete; his affective response to the idea of harming people was not fully translated to the idea of contributing to harming animals. Moralization requires re-categorization to happen

not (just) at an intellectual level, but at an instinctive one such that the moralized thing *feels* wrong.

The conflicted carnivore has made an intellectual moral judgement against meat, but this judgement is only moralized partially, or not at all. Conversely, moral vegetarians and vegans moralize meat, though to varying degrees – some occasionally lapse, whereas for others, meat is deeply suffused with moral wrongness to the point where it holds no temptation. Below, I turn to psychologically informed accounts of moral judgement. I argue that affect is important but not sufficient for making a moralized judgement, and that prevailing social norms will also have an important role to play.

3.4.2 *Reasoning against meat*

Philosophical arguments put forward in favour of one or another moral position tend to rely on an internalist model of morality where, through reason and logic, the philosopher can persuade his or her interlocutor to abandon their false beliefs, accept the better conclusion, adopt the associated moral position, and thereby change their behaviour. This certainly seems to be Singer's (1975) approach in *Animal Liberation*, and also the approach of many other philosophers writing on animal ethics. If they can show people how bad meat production is, they can persuade them that meat eating is wrong. If people see that meat eating is wrong, they will stop eating meat, and the world will become a better place for animals.

However, the evidence suggests that people's ability to reason themselves into a moral position *and act on it* is limited. Empirical research points to only a weak correlation between professed normative beliefs and action (Eagly and Chaiken 1993, Wicker 1969, via Bicchieri 2017, p. 10, 48). Some professed normative beliefs are consistently action-guiding while others play little or no action-guiding role. For example, Rust and Schwitzgebel (2014) found that ethics professors appear to embrace more stringent views regarding the requirements of morality than both non-ethicist professors and professors in other disciplines. However, ethicists are not statistically likely to behave better than their peers, let alone to live up to their professed beliefs. Most relevantly, ethics professors were significantly more likely to rate meat eating as morally bad: 60% compared to 45% of non-ethicist philosophers and only 19% of other academics (Rust and Schwitzgebel 2014, p. 96-97). However, when asked about their most recent evening meal, ethicists did not appear to be behaving better: 37% reported having eaten the flesh of a mammal the night before, compared with 33% of non-ethicist philosophers, and 45% of non-philosophers; the difference was marginal, and not in favour of the ethicists. Ethicists were somewhat more likely to report no meat consumption in the past week (27% of ethicists compared to 20% of

other philosophers, and only 13% of non-philosophers), but were nevertheless far from being consistent with their professed beliefs.

For reasons to be explored below, it seems that the moral judgements these ethicists reason themselves into play little or no action-guiding role in their lives. These judgements are not, for the most part, being moralized, and so these moral beliefs do not pose a practical barrier to divergent action. Being rationally persuaded by a moral argument is not sufficient for moralization.

If those who think about ethics for a living fail to be motivated by their intellectual moral judgements, it is hardly surprising that others also fail to reason their way into changing their behaviour. Indeed, this is the puzzle of the conflicted carnivore: such people seem to have reasoned their way to a conclusion without following through on it. Something more than reason alone must be required to generate moralization.

3.4.3 *The importance of affect*

Experimental findings in the psychology of moral judgement indicate that rational processes play a lesser role in everyday moral judgement than many of us would like to believe. This has led to a resurgence in the popularity of theories of moral judgement that privilege the role of affect, including neo-sentimentalism (Gibbard 1990; McDowell 1998; Wiggins 1987), social intuitionism (Haidt 2001), and sentimental rules (Nichols 2004). The role of affective and intuitive processes is also acknowledged by those in the rationalist camp. For example, while arguing for a rationalist conception of moral judgement Kennett and Fine (2009, p. 91) admit that “Nonetheless, Haidt (2001) is right to emphasize the relative rarity of conscious, controlled processing in everyday life.” I argue that irrespective of how one conceives of ‘true’ moral judgement, we cannot ignore the role that affect plays in people’s responses to moral, or potentially moral, issues. Nevertheless, while affect is important for moralization, I argue in [Section 3.4.4](#) that it cannot fully account for the process, and so in [Section 3.4.5](#) I argue for the role of rules.

A well known example of the role of affect in moral judgement comes from social psychologist Jonathan Haidt’s (2001, p. 814) incest study. Study participants were presented with a story of brother-sister incest where all possible harms had been eliminated. It was consensual; they used two forms of contraception; they only did it once and found it enjoyable but decided not to repeat the experience; and they decided to keep it to themselves, not as a shameful secret, but as one bringing them closer together. When asked about the moral acceptability of what the siblings did, most study participants immediately responded that it was wrong. However, when asked to justify this response, study participants searched for reasons but most could only

suggest possible harms already ruled out by the scenario. In the end, they had no recourse except to insist that incest just *is* wrong, and intrinsically so, despite the lack of possible harms resulting from it in this particular case.⁸

Haidt (2001, p. 817) describes these study participants as “morally dumbfounded” – “they would stutter, laugh, and express surprise at their inability to find supporting reasons, yet they would not change their initial judgements of condemnation.” He claims that the participants made an intuitive judgement of wrongness regarding the story based on the feeling of disgust most of us experience when contemplating the idea of incest. Yet they were not satisfied with condemning the incest on this basis alone – they wanted to be able to provide other reasons for their judgement, and were perturbed by their inability to do so. *Prima facie*, it seems that in such cases, the negative affective response to the thought of incest was sufficient to generate a strong moral judgement, even absent any ability on the part of the participants to defend their position, i.e., absent any rational processes.

It is noteworthy that participants were steered towards harm-based reasons that presupposed a consequentialist moral framework, and that many participants’ reasons could be accepted on a deontic framework where rules such as ‘incest just is wrong’ can be accepted without being grounded in harm (Royzman, Kim, and Leeman 2015, p.308). I will argue for the role such rules play in moralization in Section 3.4.5. Nevertheless, the disgust response clearly plays an important role in making the rule seem an important one.

The importance of affective responses can provide the beginning of an answer to this chapter’s puzzle. Because of the role of affective responses, our process of moral reason-giving when engaging in a moral disagreement does not, except very rarely, directly change minds – our own or others’. If eating meat *feels right* to a person, reasoning with them is highly unlikely to change their mind about this. A comparable situation arises if eating meat *feels wrong*. Reasoning without affect is ineffective. What reason-giving does, when it is successful, is elicit new intuitions that change the way the facts of the case are interpreted; reason giving can make new facts of the situation salient, allowing affective responses to be reinterpreted, giving rise to a different moral judgement (Haidt 2001, p. 820). However, reason-giving will only succeed if it can cause the case in question to be reinterpreted as falling under the domain of an already strongly accepted moral rule – a re-categorization of the type discussed in Section 3.4.1 – thereby enabling moralization. Otherwise, it will not lead to a change in judgement or behaviour.

Reason-giving might sometimes directly lead someone to change their behaviour, if they are strongly motivated to act consistently with

⁸ This is not to say that sophisticated answers to why the incest was wrong are not possible, only that most study participants were unable to articulate such reasons.

their beliefs. Haidt (2001, p. 819) admits that direct reflection on the moral facts of a situation can sometimes lead to a change of mind through sheer force of logic, with philosophers being likely candidates for such a change. This accords with the Humean notion that reason can move us to act if it can reveal to us how to achieve some already existing desire, such as a wish to behave morally, whatever that may require (Hume 1739, T 2.3.3). However, direct reflection will rarely be effective, and can only succeed where intuitions about the case are weak to begin with.

In cases where intuitions are stronger, reasoning that diverges from one's affective response will lead to a *dual attitude*, where the intuitive judgement sits below the surface of the explicitly endorsed principle. The dual attitude stems from a disconnect between System I and System II processing, as discussed in Section 3.4. Where deliberative System II thinking contradicts the answer given by the intuitive System I – where the 'gut' contradicts the head – the answer given by System I will continue to feel right, even if one intellectually knows that it isn't (Greene 2009, Haidt 2001, p. 819, 2000). Experientially, the dual attitude is like viewing an optical illusion – knowing about an illusion does not make the sense of the illusion go away, and seeing past the illusion, if it is possible at all, takes effort.

Thus one might, as I did, reason oneself into the position that the incestuous siblings did nothing wrong, while nevertheless feeling uneasy about this conclusion; the reasoned System II is unable to override the affect-informed System I judgement. This is analogous, albeit in reverse, to the experience of some readers of a text like *Animal Liberation*. Such a reader might be intellectually persuaded of the wrongness of meat eating without being motivated to change. Just as Haidt's incest study participants said that the incest *just was* wrong without being able to justify why, many meat eaters feel just as strongly that eating meat *just is* permissible, even if they cannot justify their position well (as seen in Chapter 2), or if like the conflicted carnivore, they explicitly cannot justify it.

The feeling that eating meat is permissible, despite an intellectual judgement to the contrary, drives a dual attitude. System II says "wrong"; System I says "delicious". This is a key part of why arguments like Singer's so commonly fail to effect change. Even when such arguments manage to persuade intellectually, they usually fail to produce the affective response necessary for a unified judgement. Rational argument alone will rarely, if ever, be successful in motivating change – it may lead to an intellectual moral judgement, but not a moralized judgement. At most, rational argument might lead to a dual attitude where the intuitive feeling that meat eating is acceptable will continue to exert a strong countervailing force against this intellectual judgement. To borrow Haidt's (2001) analogy, trying to change someone's mind through reasoning alone is like wagging a

dog's tail to make it happy. Dogs wag their tails because they are happy, but as philosophers should know, $a \implies b$ does not mean that $b \implies a$; wagging a dog's tail will not make it happy. Likewise, people cannot be reasoned into a moralized judgement, affective input is also required. Moralization requires persuading not System II, but System I.

Thus an affective response tends to be necessary for moralization, and its role can be decisive in turning people vegetarian. Many people report their decision to become vegetarian (even if their vegetarianism was short-lived) being precipitated by an emotional bond with an animal of a species commonly used for food. Similarly, seeing distressing footage of slaughter can lead to disgust at the thought of eating animal flesh (e.g. Haidt (2006, p. 165), Hamad (2013)). It is also common to hear or read an account of someone raising an animal with the intention of killing and eating it, only to find themselves unable to follow through on the killing, or, having killed the animal finding themselves unable to enjoy the meat. Affective responses to animal suffering or death can, as Haidt suggests, lead to the moralization of meat.

3.4.4 *The problem with an affective explanation*

While affect can lead to the moralization of meat, it need not, and so the presence or absence of an affective response cannot solve the puzzle of the conflicted carnivore. Most people, confronted with and moved by the same affective stimulus that motivates the vegetarian, are not lead to moralize meat. For example, more often than not, emotional difficulty killing animals oneself will turn someone off backyard animal raising without turning them off eating meat. Similarly, many people find footage of slaughter upsetting without this causing them to disapprove of the slaughter itself – indeed, in many cases, the vulnerability to upset is interpreted as a personal failing, leading to the conclusion that one should ‘toughen up’, not that one should change one’s diet.

Avoiding negative feelings from aversive stimuli can generally be achieved either by treating the source of the stimulus (e.g. helping someone in pain), or by removing oneself from the stimulus (e.g. looking away or leaving the room). Experimental evidence suggests that in core moral cases, most people will help even when avoidance is easier, but will otherwise choose the path of least resistance (Nichols 2004, ch. 2). Hence not all instances of suffering generate a moralized desire to help.

Where negative affective responses are reasoned as being inappropriate, they can be suppressed or overridden. For example, Wheatley and Haidt (2005) found that study participants made more severe moral judgements when disgust was induced and relevant to

the scenario at hand, but were largely able to disregard induced disgust where feelings of disgust were inappropriate to the case. Negative affective responses can be ignored. In the case of meat, people have a preference not to be confronted with the distress of witnessing animal suffering, but this preference can be satisfied just as well, and achieved more easily, by avoiding and suppressing the stimulus rather than by acting to reduce that suffering. Avoiding animal suffering is, for most people, a preference; preventing animal suffering is not a value. Insofar as avoiding exposure to animal suffering is an acceptable resolution to the problem, animal suffering is not being treated as properly moral.

Negative affect relating to meat tends not to have a lasting effect. For example, after discussing how reading *Animal Liberation* failed to move him to change his behaviour, Haidt provides a contrasting description of the apparent success of affective impetus:

I watched in horror as cows, moving down a dripping disassembly line, were bludgeoned, hooked, and sliced up [...] For days afterward, the sight of red meat made me queasy. My visceral feelings now matched the beliefs Singer had given me [...] I became a vegetarian. *For about three weeks.* (Haidt 2006, p. 166, emphasis added)

Here, the visceral affective input eliminated the desire for meat altogether, *but only temporarily*. For a time, the sight of meat held a different valence for Haidt. Rather than eliciting thoughts of a delicious flavour and memories of meals enjoyed in good company, the meat instead served as a reminder of a horrific process. Hence disgust, rather than desire, was elicited. Yet this change lasted only three weeks. Given his negative affective response to meat, Haidt developed a strong preference not to eat it. However, this preference was not transformed into a value; it was not moralized. Hence, as his visceral affective response to meat faded, so did his motivation to remain vegetarian. The process of meat production was seen as *bad*, but not as *wrong* in any deep sense.

Negative affective stimulus can be seen as bad, even upsettingly bad, without thereby being seen as having wrong-making properties. Signs of animal suffering, thoughts of animal death, or disgust at animal flesh or blood, all of which can arouse disgust (among other emotions), tend to have a fleeting impact. Absent moralization, the bad thing will not be seen as demanding a response, and so avoidance will be seen as an appropriate means of dealing with the negative stimulus. An affective response is necessary for moralization. It is not, however, sufficient.

Of course, there is a normative question about the extent to which affective responses such as disgust *should* play a role in moral deliberation, i.e. whether disgust or upset ought to be treated as markers

of wrongness.⁹ I do not here take a position on the weight we ought to assign to emotion when deliberating on moral issues. Rather, the key empirical observation is that disgust and other emotions play a strong role in moralizing an issue,¹⁰ but that they do so only under certain circumstances. Given the empirical role disgust tends to play, the difference between those cases where disgust leads to moralization and those where it does not will provide a piece to the puzzle of the conflicted carnivore. Exposure to suffering need not generate a perceived moral obligation to remedy its source; it may just lead to a desire to avoid being further exposed. The affective explanation is therefore insufficient – it does not tell us whether a negative stimulus such as information about animal suffering will be treated as something unpleasant, to be avoided in whatever way is easiest, or as something wrong, to be responded to even at some personal cost.

3.4.5 *Contextualising affect: the role of convention*

The under-determinacy of the affective explanation can be remedied by understanding moralization as requiring both an affective response *and* a normative framework prohibiting the thing that caused the affective response. According to Nichols (2004), rules governing behaviour can be either moral or conventional. Moralization accompanies moral rules, but not conventional rules.¹¹ Meat eating, however, is prohibited by neither moral nor conventional rules.

According to Nichols (2004, p. 26), moral rules are seen as very serious, generalizable (they are seen to apply to (almost) everyone), and possessing extra-conventional justification (they are wrong for reasons other than just because ‘we just don’t do that’, often to do with causing harm or being disgusting). Because of this, no authority figure can give one permission to break a moral rule. For example, murder is wrong, no matter who gives you permission to commit it.¹²

Conversely, conventional rules need not be obligatory, are recognised as not being binding on out-group individuals, and are authority-contingent. For example, while we expect to shake hands in business environments, we understand that people from other cultures may have different and equally valid conventions, such as bowing. Furthermore, given permission by an authority figure to act contrary to a conventional rule, the force of the rule dissipates

⁹ See Kelly (2011, ch. 5) for an overview of the literature and an argument against relying on disgust.

¹⁰ A point with which Kelly (2011, ch. 4) agrees

¹¹ Note that convention is used here in the colloquial sense, not the strict Lewisian sense (see (Lewis 2002)).

¹² Not every case of killing will be understood as murder. Nevertheless, when killing is sanctioned by an authority figure, as in an execution or in war, an important part of the process is that it is framed as not-murder – i.e. outside the scope of the rule.

(without necessarily disappearing entirely). The meanings of conventional rules “are defined by the constituted system in which they are embedded” (Turiel, Killen and Helwig 1987 via Nichols 2004, p. 5); the meanings of moral rules are (more or less) absolute.

Conventional rules are the society-specific rules a child learns growing up – they form a normative framework governing what is forbidden, what is permissible, and what is obligatory (Nichols 2004, p. 5-8). These rules regard issues such as the correct forms of address towards other people, correct table etiquette, or the appropriate attire for a given situation. Moral rules overlap with conventional rules: like conventional rules, they are taught to children and agreed upon by society. However, moral rules have an important additional factor: in addition to fitting into this socially sanctioned normative framework, they are also backed by an affective response. Thus one experiences a visceral response to a violation, or even the idea of a violation (Nichols 2004, p. 26).¹³ While the breach of a social convention, for example the wearing of too informal attire to an event, does not usually provoke a strong response in an observer, a breach of a moral rule does. Because of the affective component, moral norms are held more strongly than conventional norms – the very idea of violating them tends to be upsetting in a way that the idea of violating a conventional norm is not.¹⁴ While both moral and conventional rules have normative force, the normative force of moral norms transcends mere convention, while still relying on convention.

Without a rule, an upsetting thing is neither a moral nor a conventional violation. As a society, we do not have rules against killing or eating animals, and so meat eating and associated processes are not treated as moral issues even when they elicit a strong negative affective response. Without the rule component, the affective component can be stifled. This explains why many issues that generate an affective response, including meat, are not treated as moral – rightly, in many cases. Many things that generate an aversive response are not, and should not be, treated as immoral. For example, vaccinating a reluctant child is the right thing to do, even if the child’s protestations are upsetting. Similarly, seeing someone punched in a boxing match elicits an affective response, but we do not normally think their opponent did anything wrong. Such behaviour is expected and consented to in that context, and so disapproval tends to be stifled even if one personally objects to boxing. Likewise, while the sight of

¹³ The affective response need not be present in every case, but it is likely that if a person failed to experience the affective response for a prolonged period of time, the rule would come to lose much of its force (Nichols 2004, p. 27).

¹⁴ The distinction between moral and conventional rules is not, however, always clear cut. Some rules are treated as being moral by some but conventional by others, for example, the permissibility of smoking in a school, or the permissibility of eating pork. Hence some people may respond viscerally to what seems to someone else a merely conventional violation.

blood can elicit disgust, we do not treat someone who is bleeding as having done something morally wrong – we do not have norms against bleeding because we recognize that this is something outside a person’s control. Absent the normative component, an upsetting situation is not treated as wrong.¹⁵ Hence a negative affective response alone, like that experienced by Haidt, is not enough to motivate a moralized response to meat or any other issue.

Furthermore, social pressures mediate affective responses. For example, while some things appear to be innately and universally disgusting, such as “the pungent smell of rotting garbage on a hot summer day”, many disgust responses are learned, and what is seen as disgusting can vary across individuals and cultures (Kelly 2011, p. 11). For example, the idea of eating pork may be disgusting to a Jew or Muslim while being appetizing to people from other backgrounds. Social pressures thus play a role in people’s affective responses to meat-related stimuli.

For the most part, people are normatively encouraged to tolerate or avoid the aspects of meat that have the potential to generate negative affective responses. For example, raw meat frequently generates a disgust response in children and some adults (especially those who have not been acclimatised to handling it), but the normative framework in which most of us are raised encourages desensitisation (i.e. learning to suppress the disgust response in order to handle raw meat) rather than moralization. Similarly, undercover videos taken in factory farms and slaughterhouses are frequently highly distressing, triggering our affective aversion to seeing deliberately inflicted harm. Yet as Melanie Joy (2010, p. 132) notes (and as seems to have happened in Haidt’s case), this distress usually wears off with little or no long term effect. Furthermore, in most cases, even this temporary upset can be avoided simply by refusing to look. Most societies’ normative frameworks have generic rules prohibiting causing unnecessary harm, but killing animals for food is seen as necessary,¹⁶ and so our normative framework encourages us to turn away rather than to moralize. Lacking normative reinforcement, information about the problems of meat production and its attendant negative feelings are readily forgotten. Thus meat eating is seen as morally acceptable, or even desirable, despite the negative affective responses it can engender.

Not only does the absence of normative rules against something such as meat cause it not to be treated as moral, convention also im-

¹⁵ Many vegetarians and vegans treat meat eating as wrong, despite the lack of normative rules discussed. This can be explained as they employment of an idiosyncratic understanding of the rules extant in their society, for example, by treating rules against causing harm as applying to a broader than usual target group (including animals) (Nichols 2004, p. 156).

¹⁶ The necessity of eating meat is frequently put forward as a justification for the moral acceptability of killing animals (Bastian *et al.* 2012), as was discussed in Chapter 2.

portantly mediates who is, and who is not, a possible subject of moral violation. In such cases, an affective response to a stimulus will be present (to varying degrees) in most people. However, although the affective response is present and the act is of a sort otherwise prohibited by convention, the target of the act is such that the conventional prohibition either is not triggered or is overridden by a norm of explicit approval. Our society has rules against killing, and people experience an aversive response to witnessing it, but these conventional rules do not extend universally. They apply to human beings and are extended to some animals, particularly pets, but they extend neither to animals used for food nor to those considered “pests”.

Being outside the group to which social norms forbidding harmful acts apply can account for how we interact with animals of species commonly killed for meat, as well as for our historical treatment of other out-groups. Nearly every culture studied by anthropologists has norms prohibiting a range of harmful acts (Nichols 2004, p. 142). However, there is significant variation both in what is considered harmful and whom it is impermissible to harm (Nichols 2004, p. 142-143). In no culture is all killing prohibited, though a few cultural groups, Jains among them, come close.

Norms determine which killings are wrong and which are outside the scope of morality. In every culture, some plants can be killed; in almost every culture, some animals can be killed; and in most cultures, some people can be killed. Today, the killing of farmed animals is treated as amoral, but historically, many peoples were also excluded from the scope of norms against killing. For example, the attempt of British settlers to exterminate the Australian Aborigines was a case where the killing of humans was viewed in an amoral way. The killing of Aboriginal people was, appallingly, seen as morally good insofar as it improved the lot of white Australians, as the lives of Aboriginal people themselves were not considered to have moral value (Ryan *et al.* 2017). Aboriginal people were therefore not covered by then-existing norms against killing. In the historical case, an affective response would have been experienced by most on witnessing the death of any human, but as the norm against killing only extended to cover white settlers, the killing of Aboriginal Australians was, for many, not moralized. Today, in Western societies, the norm against killing extends to all people (except, to some extent, criminals and enemy combatants), but does not, for the most part, extend to animals. Thus in Western cultures such as Australia, the killing of humans and some non-human animals (pets and endangered or cute species) is moralized, while the killing of other animals, particularly animals regularly used for meat, is not moralized. It is possible for individuals to extend harm norms to cover animals – this is what vegetarians and vegans who moralize meat achieve, and is also how our current normative bounds have been stretched from more circum-

scribed historical bounds (Nichols 2004, p. 156). However, doing this is not straightforward.¹⁷

Not only do our norms explicitly exclude ‘meat’ animals, there are other norms endorsing meat eating. Thus while the affective factor separates moral rules from conventional rules, convention also circumscribes morality. It determines who is a member of the moral community – who is protected by the rules and who is not (Nichols 2004, p. 142-143). Convention mediates the affective response, telling one whether or not a given target is an appropriate locus of moral concern. Because of prevalent social rules, negative affective responses to meat-related prompts tend to be stifled rather than treated as indicating a moral issue.

3.5 CONCLUSION

Many people eat meat despite claiming to think meat eating morally wrong. Such people, conflicted carnivores, manifest an interesting disconnect between their professed moral judgements and their behaviour. This disconnect cannot be adequately explained by calling the moral judgement insincere, by rejecting motivational judgement internalism, or by calling the conflicted carnivore weak-willed. Each of these explanations, even if accepted, raises deeper questions about the difference between the conflicted carnivore’s relationship to meat and his or her relationship to other moral issues where motivation to follow through on moral judgement is present. Rather, understanding this disconnect begins with observing that irrespective of whether or not the conflicted carnivore’s judgement should rightly be considered a genuine moral judgement, it is not a moralized judgement. For the conflicted carnivore, eating meat is judged to be wrong only in the abstract; avoiding it is not an internalized value. This contrasts with the attitude of many committed vegetarians, for whom the perceived immorality of meat makes it unappealing. It also contrast with the conflicted carnivore’s own behaviour in other parts of life, where doing the perceived-to-be-immoral thing would be unthinkable. A good explanation of the conflicted carnivore’s motivational deficit needs to address the basis of this difference, which requires understanding what is involved in making a moralized judgement.

While simply thinking about the issue *may*, in rare cases, be enough to motivate behavioural change, I argued that affective input is an essential ingredient for making a moralized judgement. Much of the conflicted carnivore’s behaviour can therefore be understood by observing that although their intellectual System II reasoning says that eating meat is wrong, the intuitive, affect-driven System I says that meat is morally fine. In cases of such a dual attitude, resisting the verdict of System I is effortful and therefore unlikely to be sustained.

¹⁷ I explore this idea further in [Chapter 4](#).

However, such an explanation is still insufficient. In many cases, such as after viewing graphic footage of slaughter, meat can be accompanied by significant negative affect without this generating moralization. In such cases, System I and System II can be in agreement, at least for a time, without this generating sustained change.

This affective under-determinacy can be remedied by recognizing that moralization requires both an affective response to something, and a normative framework prohibiting such actions. This approach can make sense of the fact that often the response to an unpleasant stimulus is not a moralized response, but rather a desire to simply avoid it – absent a normative framework, the stimulus indicates something *bad* without indicating something *wrong*. Thus the conflicted carnivore who has been exposed to a meat-related affectively upsetting stimulus may turn away; our culture's normative framework explicitly tells us that such things are not a sign that something wrong is happening. Hence whatever affective response a person might feel when exposed to the reality of animal suffering for meat production is likely to be attenuated; the reality of the violence of meat is normatively de-emphasised. Furthermore, the location of the harm is outside the standard socially recognised normative framework – we do not treat farmed animals as being the appropriate targets of norms against causing harm, and so any affective response is likely to be processed as aberrant. It is not treated as a reliable indicator of a moral issue, but is instead seen as something which can and should be dismissed.

Nevertheless, some people do internalize moral rules which do not have widespread acceptance, vegetarians and vegans being two such groups. In order to better understand this phenomenon, it is necessary to turn to an analysis of how norms surrounding meat eating have formed, how they currently operate, and how they might be changed. In the next chapter, I turn to a more rigorous explanation of the development and contemporary function of these norms, and why some people but not others are able to challenge them successfully.

NORMS OF MEAT

4.1 MEAT EATING AS A COLLECTIVE BEHAVIOUR

In our society, meat eating is the default. It is therefore easy for people to avoid seeing meat as a moral issue (Chapter 2) and to avoid treating it as a moral issue even if they are intellectually convinced that they should (Chapter 3). In this chapter, I turn to an analysis of how current social dynamics support meat eating practices and what this says about the factors that could precipitate change. Following Christina Bicchieri's terminology, I argue that meat eating is largely a *descriptive norm* but also has elements of a *social norm*. Because of this, approaches that focus on directly persuading people to change their behaviour are unlikely to be effective and may even backfire. Rather, a norms-based approach that recognizes that people are sensitive to both normative and empirical expectations is more likely to influence change. I then argue that if enough people could be motivated to change, a *tipping point* could be reached where anti-meat norms take hold and, over time, the eating of animals would become moralized.

I begin, in Section 4.2, by following Pettit and Brennan (2005) in suggesting that an important motivator for moral behaviour is a concern for seeking esteem and avoiding disesteem (what Adam Smith (1774) calls "approbation"). Hence people will adjust their behaviour to increase expected esteem. Social norms can therefore incentivise moral behaviour, but only where people expect to receive esteem for complying with norms or disesteem for violating them. Because meat eating is so widespread, it is close to neutral with respect to esteem, and so in our current social environment, esteem considerations are unlikely to motivate or sustain meat-avoiding behaviour.

Thus understanding how people think about the morality of meat eating and the role esteem considerations play requires understanding the nature of this behavioural regularity. Therefore, in Section 4.3, I turn to an analysis of the practice of meat eating, drawing on Cristina Bicchieri's (2006, 2017) recent work on social norms and collective practices. Bicchieri divides behavioural regularities into customs, descriptive norms, and social norms. *Customs* arise where people independently converge on a behaviour because it is individually advantageous (Bicchieri 2017, p. 15). *Descriptive norms* arise where people converge on a behaviour because converging is rational given what others are doing (Bicchieri 2017, p. 19). Finally, *social norms* arise where people converge on a behaviour because of both first order empirical expectations and second order normative

expectations (Bicchieri 2017, p. 35). I argue that meat eating does not straightforwardly fit into any of these categories, as different people within Western societies conform for different reasons. Nevertheless, I argue that meat eating as a custom is relatively rare and that for most people, it is a descriptive or social norm. Thus any useful analysis of meat eating must take into consideration the way these norms shape behaviour.

I then, in Section 4.4, turn to an exploration of what it would take to change existing practices surrounding meat consumption. I argue that while some people have a strong independent drive to eat meat, most people will be sensitive to various incentives. Because meat eating is widespread, conforming to the meat eating behavioural regularity is an easy way to meet nutritional needs. However, if fewer people ate meat, it would be less accessible and possibly also more costly. Hence without the descriptive norm of meat eating in place, meat eating would be less easy and therefore less attractive. I then note that there are also people for whom the social norm aspects of meat eating have a strong influence – such people would be open to change absent social pressures to conform to meat eating. Hence small changes in the second order normative expectations of such people could result in behavioural changes.

It is therefore important to understand the factors that enable or impede norm change. That is the focus of Section 4.5. *Norm sensitivity* and *risk perception* both influence whether someone will be prepared to deviate from an established social norm (Bicchieri 2017, p. 172-182). Education campaigns can undermine norm sensitivity by undermining reasons for meat eating and providing reasons not to eat meat. However, reducing norm sensitivity is necessary but not sufficient for norm change; such approaches ignore the role of risk. A potential vegetarian also needs to be prepared to break with tradition – with socially expected behaviour – which requires that they perceive the risks of transgressing meat norms to be low enough to warrant action (Bicchieri 2017, p. 166).

I conclude by arguing that if anti-meat norms were to become established, then the risk of sanctions, such as disesteem for eating meat, would motivate compliance in most current meat eaters, and would, over time, lead to the moralization of meat eating. Hence, if one thinks that as a society we should eat less or no meat, my analysis indicates a possible route for widespread change.

4.2 MEAT AS SOCIALLY ENDORSED

One's social environment influences both how one thinks about meat and one's willingness to reduce or eliminate meat consumption. Many people think that for moral reasons most people should eat less meat – 50% of respondents in our survey agreed (Thrasher

and Mayne 2018). Yet given prevailing norms, most of these people continue to eat meat. Many others, or so I argue, would conform to predominant vegetarianism, as people do in parts of India where lacto-vegetarianism is the norm. Even if this conformity to vegetarianism were initially reluctant, I argue that given time and a conducive normative environment, many initially reluctant conformers would come to endorse anti-meat norms and treat meat eating as immoral. However, since the vast majority of people interact predominantly with meat eaters, they have no social motive to change. Furthermore, many people, while sympathetic to arguments for vegetarianism, express concern that their independently going vegetarian would lead to little or no improvement in outcomes for animals (e.g. Budolfson (2016)). Hence both willingness to reduce or eliminate meat consumption and the ways people think about the morality of meat eating are significantly dependent on what other people do.

Because of meat eating's prevalence, it tends not to seem particularly bad, even for people who agree on an intellectual level that eating meat is immoral. As explored in Chapter 3, even people who agree that eating meat is morally wrong continue, for the most part, to eat meat with little guilt. Furthermore, even most strict vegans do not properly treat meat eating as a moral issue. For example, most moral vegetarians and vegans would avoid associating with an unapologetic racist but will happily associate with unapologetic meat eaters. This inconsistency in how meat is treated compared to other moral issues is, I argue, directly tied to just how common it is.

The subjective urgency of complying with a moral demand tends to vary depending on how prevalent compliance is. Many studies find only a weak correlation between personal normative beliefs and action (e.g. Eagly and Chaiken (1993) and Wicker (1969)). Stated personal normative beliefs are insufficient to determine action. Nor do they seem to have much influence. Rather, where empirical expectations about what others will do diverge from personal normative beliefs, it is the empirical expectations which tend to drive behaviour (Bicchieri and Xiao 2009). Thus, given that most people eat meat, individual normative beliefs about the immorality of meat, even where present, are unlikely to motivate vegetarianism on their own.

People care about their own self-interest but also care about behaving (or being perceived as behaving) appropriately. As Adam Smith noted, on average, people's behaviour can be best explained by assuming them to be fundamentally self-interested:

It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity, but to their self-love, and never talk to them of our own necessities, but of their advantages. Nobody but

a beggar chooses to depend chiefly upon the benevolence of his fellow-citizens. (Smith 1776, p. 8)

Yet Smith also notes that self-interest does not preclude people from valuing morality and behaving in self-sacrificing ways:

How selfish soever man may be supposed, there are evidently some principles in his nature, which interest him in the fortune of others, and render their happiness necessary to him, though he derives nothing from it except the pleasure of seeing it. (Smith 1774, p. 11)

In evaluating a course of action, people care both about the overall outcomes likely to be produced by that course of action and the personal costs and benefits. For someone who thinks that animals matter morally and that eating meat contributes to unjustified animal suffering, avoiding (intensively farmed) meat leads to good outcomes, all things considered. However, these all-things-considered goods come at a personal cost, and self-interest will incline one towards choosing the lower-effort and more enjoyable option, which is, for most people, continuing to eat meat. Thus the feasibility of ‘moral’ behaviour is constrained by self-interested motives. If a course of action leads to all-things-considered good outcomes, but not personally good outcomes, then it is unlikely to be perceived as either attractive or feasible.

For someone who enjoys eating meat, reducing or eliminating meat is a self-sacrificing behaviour – it involves effort and a certain loss of enjoyment.¹ Hence for the conflicted carnivore described in [Chapter 3](#), there is a conflict between the perceived moral behaviour – vegetarianism – and what they want – continued meat consumption. However, an important self-interested motive factoring into cost-benefit analyses is the desire to gain esteem and avoid disesteem, i.e. the desire for a good reputation (Pettit and Brennan 2005; Smith 1774). This consideration has the potential to influence decisions about meat consumption, but it currently has little effect.²

Unlike with straightforwardly moral issues, there are few social factors to induce vegetarianism, making vegetarianism a relative rarity. Moral behaviour apparently counter to self-interest can largely be explained by people behaving in ways that will enhance their image.³ The level of esteem or disesteem to be gained for particular

¹ Avoiding meat need not be seen as a sacrifice, as will be explored in [Section 4.5](#). However, since most meat eaters will at least initially frame avoiding meat as a sacrifice, this perspective must be used as a starting point for understanding the practice of meat eating.

² Thanks to Michael Smith for explicitly drawing my attention to the interplay of self-interest and esteem in generating moral behaviour.

³ Even invisible moral behaviour can be explained by such an account: internalizing a value and acting on it consistently may be a more effective route to building esteem than making a calculated decision each time.

behaviours is determined by what other people are doing; the degree of esteem gained for 'good' behaviour or disesteem incurred for 'bad' behaviour is affected by the baseline for that behaviour in one's group (Pettit and Brennan 2005, p. 19). If everyone is acting well, compliance gains no esteem, but divergence attracts disesteem. Conversely, if everyone is behaving badly, bad behaviour will not attract disesteem, but good behaviour may attract esteem.

Thus one's moral ideals may matter less than what the average person around one does. Because most people in mainstream Western societies eat meat, meat eaters do not tend to attract disesteem, even from vegetarians who feel strongly about the immorality of meat eating.⁴ Given the current baseline, one can gain esteem by purchasing free range meat or following 'meat free Mondays' without needing to go all the way to vegetarianism or veganism. It should therefore be unsurprising that there are significantly more conflicted carnivores than there are vegetarians.

Furthermore, while there is some potential to gain esteem for 'doing the right thing' when few are, morally exemplary behaviour can bring disesteem rather than esteem. People tend to seek the esteem and avoid the disesteem of those in their reference network – the set of people whose opinions matter and hence to whose social norms one is inclined to conform. The reference network tends to be the people one interacts with regularly – family, friends, and perhaps colleagues – but can also be constituted primarily by geographically removed people such as members of an expatriate's home culture (Bicchieri 2017, p. 19, 33-34). Because of this, people will, other things being equal, tend to carry out acts that attract esteem and avoid acts that attract disesteem from those they care about with little thought for the esteem of people outside their reference network (Pettit and Brennan 2005, p. 19). However, which acts will attract esteem, disesteem, or be esteem-neutral varies across time and place. Different cultures value different things, but even more important than the broad culture of one's society are the values of the subculture or peer group in which one is embedded. The attitudes of mainstream society may not matter much if one is deeply embedded in a subculture with different values. Thus vegetarianism's esteem implications will vary depending on one's reference network. In most reference networks, however, vegetarianism will not bring esteem.

This can be because of a disagreement about what morality requires, or due to a moral 'tall poppy syndrome'. In the first case, meat advocates might think vegetarians are doing something wrong, for example, putting farmers' livelihoods in jeopardy or advocating for a reduction in welfare for people living in developing countries.

⁴ However ex-vegetarians seem to be judged more harshly, suggesting that they are judged against a vegetarian reference group rather than broader society.

Hence a disagreement about what morality requires can straightforwardly lead to a loss of esteem for ‘moral’ behaviours.

However, one can also attract disesteem for agreed-to-be-moral behaviour. We tend to respect and esteem people who are trustworthy, generous, or upstanding, yet feel uncomfortable around exemplars. As Susan Wolf (1982) notes, if someone follows their principles too strongly, we may respect them but find being around them unpleasant. Someone who tries too hard is often perceived as a ‘goody two shoes’ motivated by a desire for superiority rather than genuine morality. Furthermore, judging something to be morally wrong requires blaming or condemning those who carry out that action, and people who lay blame or condemnation too readily are unpleasant to be around (Baier 1965, p. 1). People tend to react negatively to others who are perceived as being more moral than themselves, particularly when they think they are being negatively judged (Minson and Monin 2012). Because most people see themselves as moral individuals, being confronted with evidence that others are behaving significantly more morally can prompt a backlash against the apparently more moral person. This backlash enables a positive self-image to be maintained without behavioural change (Rabin 1994; Rothgerber 2014). Hence attempts to persuade people to abandon meat eating can reinforce the meat eating status quo rather than undermining it, as people who are ambivalent about their meat eating may feel challenged by vegetarians and respond by pre-emptively judging the vegetarian. Vegetarianism can therefore bring disesteem rather than esteem.

Because of these factors, it does not make sense, socially, to aspire to be moral as such. Rather, it makes sense to calibrate one’s moral behaviour so as to be approximately as moral, across various issues, as one’s peers – slightly better in some respects, slightly worse in others, but aiming for a “B+” overall (Rust and Schwitzgebel 2014; Schwitzgebel and Rust 2015). From an economy of esteem perspective, this is reasonable; people calibrate their level of ‘moral’ behaviour to bring close to maximum esteem (Pettit and Brennan 2005, p. 16). Doing things that would cause reproach will bring disesteem and should be avoided (in mainstream Western society, things like stealing, lying, or eating dog-meat), but excessive goodness, especially that which might make one seem sanctimonious, should also be avoided. Compliance with what one thinks one morally ought to do will therefore be calibrated to be neither significantly more nor less self-interested than what others do. What other people within one’s reference group do therefore matters – this sets the baseline against which one’s actions will be esteemed or disesteemed.

For example, as mentioned in Chapter 3, empirical research by Joshua Rust and Eric Schwitzgebel (2014; 2015) found that ethics professors tend to have much more stringent views than other profes-

sors about what morality requires, but are not statistically more likely to behave more ethically as a result of those views, let alone to live up to their own standards. Of particular interest is the finding that ethics professors were significantly more likely to rate meat eating as morally bad, but when asked about their most recent evening meal, did not appear to be behaving better than other groups. In Western societies, and in most subcultures within Western societies, behaving in line with one's peers enables one to be "moral enough" while still consuming meat. After all, our empirical expectations are that most people – including most people who say meat is bad – will eat meat. It is therefore unsurprising that the 'moral' behaviour of avoiding meat is not prioritized.

Because meat eating is so widespread, disesteem for engaging in the practice is unlikely. In mainstream Western society, eating meat is close to neutral with respect to esteem. In some circles, conspicuous consumption of meat can raise esteem, particularly when it is used to signal masculinity (Rothgerber 2012; Ruby and Heine 2011). In other circles, avoiding meat can bring esteem for environmental, health, or moral reasons. However, most people do not feel strongly about the diets of others, though some meat eaters may feel antipathy towards conspicuous moral vegetarianism (Minson and Monin 2012). Given how common meat eating is, it therefore makes little sense, in the general case, to disesteem people for eating meat. One can still earn a "B+" in morality while eating meat.

Furthermore, the knowledge that the 'immoral' option is the norm increases the likelihood that one will also take the immoral course of action – and in mainstream Western society, everyone knows that just about everyone eats meat. Knowledge that the immoral action is commonplace reduces the risk of sanctions or disesteem. Yet many moralizing campaigns emphasise the prevalence of the behaviour they wish to eradicate. For example, many animal activists highlight the number of animals killed each year for food, an approach that indirectly reinforces the fact that meat eating is incredibly widespread. As Cialdini *et al.* note:

Within the statement "Look at all the people who are doing this undesirable thing" lurks the powerful and undercutting normative message "Look at all the people who are doing it." (Cialdini *et al.* 2006, p. 5)

The influence of normalizing an undesired practice is backed up by diverse experimental studies. For example, participants in a dictator game, where one person (the dictator) is given a sum of money and is free to decide how to distribute that sum between themselves and another 'player', were more likely to keep more for themselves if told that other participants had also kept more for themselves (Bicchieri and Xiao 2009). Similarly, above-average theft was observed when visitors to Arizona's Petrified Forest National Park were told "Your her-

itage is being vandalised every day by theft losses of petrified wood of 14 tons a year, mostly a small piece at a time" (Cialdini *et al.* 2006, p. 5). Similarly, the fact that meat eating is so widespread makes it difficult to combat. Not only does its widespread nature indicate that one will not be disesteemed for eating meat, it indicates that many others are reaping the benefits of this behaviour such that individuals who abstain are unreasonably sacrificing their self-interest. The problem of meat is therefore not a problem of changing the behaviour of individuals, but of changing whole groups.

4.3 ANALYSING THE COLLECTIVE BEHAVIOUR OF MEAT EATING

Most people around the world eat meat regularly, or would if they could afford to.⁵ Eating meat is therefore both statistically the norm and seen as normal; following a vegetarian or vegan diet is neither. For anyone interested in changing meat-eating behaviour, it is therefore important to understand the factors driving meat eating – in particular, the extent to which those factors are individual and the extent to which they are social. Many philosophers (including Bicchieri (2006, 2017), Brennan *et al.* (2013), Lewis (2002), Pettit (2002), and Vanderschraaf (1995)) have provided accounts of norms and collective behaviours. This section provides an overview of the different types of collective behaviours and argues that meat eating is shaped by both empirical and normative expectations.

I use Cristina Bicchieri's (2006, 2017) framework as an entry point for analysis because her account is operational – unlike other available norm frameworks, it allows both the identification of the conditions under which a norm exists and the conditions under which it is followed, and also suggests ways of identifying norms independently of observed behaviour (Bicchieri 2017, p. 65). It is also the most extensive empirically informed philosophical account not only of how to understand norms, but also of how they may change. Nevertheless, despite important differences, her account is largely compatible with the claims of other theorists such as those listed above, and for the analysis that follows, the differences between these frameworks are unlikely to have much, if any, effect.

⁵ Meat consumption rises with wealth in developing countries. As incomes rise above poverty levels, people in most countries are eating more meat than ever before (Murphy and Allen 2003).

4.3.1 *Types shared behaviours*

On Bicchieri's (2006, 2017) system of classification, a regularity in collective behaviour can be a custom, a descriptive norm, or a social norm.⁶

4.3.1.1 *Customs*

The simplest form of shared behaviour is a custom. A custom is a "pattern of behaviour such that individuals (unconditionally) prefer to conform to it because it meets their needs" (Bicchieri 2017, p. 15). Social expectations play no role in the rise of a custom (Bicchieri 2017, p. 16). People just happen to converge on a behaviour because it individually meets their needs, irrespective of what those others are doing. For example, umbrella use can arise as a custom because people find it useful to protect themselves from the rain. Whether others are also using umbrellas is irrelevant. Because adhering to a custom enables people to meet a need, individuals are likely to persist in the customary behaviour regardless of what they observe (Bicchieri 2017, p. 15). Thus, for example, umbrella use during rain is likely to persist irrespective of what other people are doing. Regularities of this form therefore result from individual rather than social preferences. Conformity to a custom is not driven by a desire to conform, but is instead a side effect of people's individually rational behaviours.

4.3.1.2 *Descriptive norms*

A descriptive norm is a regularity that results from conformity to empirical but not normative expectations. It is "a pattern of behaviour such that individuals prefer to conform to it on condition that they believe that most people in their reference network conform to it (*empirical expectation*)" (Bicchieri 2017, p. 19, emphasis original). For compliance with a descriptive norm, it is enough that one expects that others will behave in a certain way – given that other people are conforming to the regularity in question, one's best mechanism to reach one's goals is to also conform. For example, the side of the road on which one drives is a straightforward descriptive norm – the choice of side is not significant, but everyone must converge on the same one. Rules of etiquette are also descriptive norms: one must follow them to avoid being thought uncouth, but failing to conform will hurt oneself far more than others. Fads or fashions are also descriptive norms. For example, in a desire to appear fashionable, a woman might wear skirts of a certain length. Given that (almost) all fashionable women are currently doing this, if she wants to be considered fashionable,

⁶ Bicchieri's definitions of these terms differ from both ordinary usage and usage by other philosophers. Thus the definitions given below should be borne in mind when considering my later analysis.

she should do the same. Yet at least *prima facie*, there is no expectation that she should conform if being fashionable is not something to which she aspires. Descriptive norms motivate compliance because not conforming with the descriptive norm directly makes it harder for one to achieve particular goals.

Descriptive norms can be actively harmful. For example, bribery is a descriptive norm that leads to worse overall outcomes in societies where that norm is followed. Nevertheless, if one has the empirical expectation that bribery is common in a society, then it is prudent to take and offer bribes, even if one disapproves. While bribery can take on aspects of a social norm (for example, those who take a stand against bribery can be punished by those who strongly benefit from corruption), it can be self-sustaining even without any normative expectations. Conforming to the descriptive norm enables people to meet their goals – getting something done for the briber, and getting a bit of extra cash for the bribed. The empirical expectation that bribery is common means that it neither makes sense for someone to refuse to bribe (they should expect their task will not be done or will take far too long), nor to refuse to take a bribe (they are missing out on an income stream they could enjoy without any likely sanctions). Taking the moral high ground will be costly – deviating from a descriptive norm, although it can hurt others as well, primarily hurts the deviator. The expectation that bribery should (pragmatically) take place therefore exists alongside the normative belief that bribery should not occur. Hence descriptive norms can exist without supporting normative expectations, and can even run counter to them.

4.3.1.3 *Social Norms*

The last type of collective behaviour is a social norm. Social norms are maintained through both empirical and normative expectations. They usually solve collective action problems where the individually rational action and the collectively rational action pull in different directions. Through the weight of normative expectations, sometimes strengthened by sanctions, social norms make it individually rational for people to conform to the collectively rational option:

A social norm is a rule of behavior such that individuals prefer to conform to it on condition that they believe that (a) most people in their reference network conform to it (*empirical expectation*), and (b) that most people in their reference network believe they ought to conform to it (*normative expectation*). (Bicchieri 2017, p. 35, emphasis original)

For a social norm, one not only expects that others will comply with the regularity, one also expects that others expect that one should comply (a second order normative expectation). For example, one

might comply with a norm of reciprocity not only because one observes others reciprocating, but because one expects disapproval otherwise. This may further be supported by personal moral or prudential beliefs, e.g., that it would be wrong or imprudent not to reciprocate, but this is not necessary (Bicchieri 2017, p. 126).

Unlike with descriptive norms, where coordination is rational simply in virtue of being the best way of reaching a goal, the rationality of following a descriptive norm comes from the costs imposed on defectors. These costs can range from mere disapproval to active punishment. Aversion to these costs is what, in the general case, overcomes the temptation to defect (i.e. to act in a straightforwardly self-interested way). A social norm can become internalized such that compliance becomes valued for its own sake even in the absence of others' normative expectations (Brennan *et al.* 2013, p. 196). However, because there are often costs to compliance when others defect, compliance with a social norm tends to break down in the presence of significant defection, even among those who have internalized the norm (Bicchieri 2017, p. 75, 153-154). Hence the possibility of sanctions plays an important role in maintaining social norms.

4.3.2 *Categorizing meat eating*

There is undoubtedly a behavioural regularity surrounding meat eating: most people eat meat. However, determining the form this regularity takes is not trivial, and is complicated by the fact that, as Bicchieri (2017, p. 3) notes, the same practice can take the form of a custom in one population, a descriptive norm in another, and a social norm in yet another. Furthermore, within one population, people can conform to a regularity for different reasons, both at the subgroup and the individual level (Bicchieri 2006, p. 204-205). Such complications apply to meat eating.

At first glance, meat eating might seem to be a custom. It is not, however, a straightforwardly independent behaviour. In the basic case, meat eaters desire meat because they enjoy the taste and because it allows them to meet nutritional goals. Hence, meat eating does not seem to be driven by conformity to either the normative or empirical expectations of others. However, classifying meat eating purely as a custom fails to explain important facets of the practice. In any given society, the potentially nutritious flesh of many species of animal, eaten in other societies, is not only not eaten, but not even recognized as food. Individual preference cannot explain why, in any given society, people tend to eat the meat of only a few animals, with those few varying from society to society. Dogs, pigs, cows, whales, and horses are all eaten in some parts of the world while their meat is viewed with revulsion in others. This strong and moralized distinction between 'edible' and 'non-edible' animals indicates the

presence of internalized norms rather than people merely acting on individual preferences.

Furthermore, the ease of using meat to fulfil one's individual needs is a function of meat's ready availability, which in turn relies on meat eating being widespread. Empirical expectations about what others will do are therefore strongly tied to empirical expectations about the ease of sourcing meat. Thus while meat eating may simply be a straightforward way of meeting a need for some (in particular, for those who will happily eat the flesh of any animal), or a personal norm for others (particularly those who would hunt or raise their own animals even if meat were not readily available), it is not, for the most part, simply a behaviour stemming from individual preferences. Expectations about what others will do play a role in both what one can obtain and what one is willing to countenance eating.

The next option is a descriptive norm. Most societies, and all Western societies, have converged on meat-containing meals as a solution to our nutritional needs, and it can therefore be difficult to obtain tasty and nutritious meat-free meals, especially outside urban centres. Often the only options might be chips or a garden salad – an unsatisfactory meal. The empirical expectation that others will eat meat leads to the valid inference that meals containing meat will be easy to obtain, whereas sourcing meat-free meals will take non-trivial effort. Furthermore, following the crowd tends to be comfortable, especially when one is uncertain as to the best path (Bicchieri 2017, p. 22). Eating meat, like following broad fashion trends, allows one to fulfil a need in a straightforward manner while not standing out. Conversely, vegetarianism makes one stand out in a way that many find difficult. A study by the Humane Research Council (2014, p. 7) found that 63% of former vegetarians disliked the way their vegetarian diet had made them “stick out from the crowd”. Thus the co-ordination point people have reached of building meals around meat makes deviation from the norm costly, both socially, in terms of being perceived as different, and practically, in terms of being able to find alternatives.

The empirical expectation that others eat meat (and do not eat vegetarian meals) generates the expectation that not eating meat oneself will be difficult. Especially as most people also enjoy the taste of meat, this expectation is sufficient to ensure compliance with the norm in the absence of strongly motivating reasons for non-compliance. Because deviating from the norm of meat eating primarily hurts the deviator, both in terms of effort and enjoyment, meat eating seems a descriptive norm. Yet something that starts off as a descriptive norm can take on a normative valence, and this has happened, to an extent, with meat eating.

Meat eating involves aspects of a social norm while also being supported by other social norms, in particular norms of hospitality. Not

eating meat can complicate social interactions. For example, Michael Pollan recalls from his period of vegetarianism that:

Other people now have to accommodate me, and I find this uncomfortable [...] As a guest, if I neglect to tell my host in advance that I don't eat meat, she feels bad, and if I *do* tell her, she'll make something special for me, in which case *I'll* feel bad. (Pollan 2006, p. 314, emphasis original)

Traditional norms of being a dinner guest require one to eat what is provided and to express gratitude for it. Following a non-standard diet can inconvenience others, leaving room for social pressure to conform. Although there is now an increased awareness of and willingness to cater to dietary restrictions for health, religious or moral reasons, many people (35% in our survey) will still negatively judge a guest who fails to adhere to these traditional hospitality norms (Thrasher and Mayne 2018). Thus announcing that one will not eat meat can be socially fraught.⁷

Many people contemplating vegetarianism also report a fear of social sanctions beyond the above described discomfort, and most people who avoid meat have run into some level of hostility because of their dietary choices. Many meat eaters value the practice of meat eating and defend it against the threat to its continuance posed by meat abstainers. Thus a non-trivial number of meat eaters will go out of their way to mock vegetarians, treat them with hostility, and pressure them to eat meat. In doing so, they enforce norms of meat eating, even if only through relatively mild social sanctions. Although these sanctions are generally mild, a fear disapproval or mockery can be sufficient to dissuade someone who is only weakly motivated. Furthermore, men who avoid meat are rated as less masculine (Rothgerber 2012; Ruby and Heine 2011), and thus normative expectations regarding how a man *should* behave can reinforce meat-eating behaviour in men, adding an extra challenge for men considering vegetarianism. The fact that meat eating can be normative is particularly obvious in moral vegetarians who reluctantly eat meat to keep the peace in various social situations. My brother did this around my father for over a year! Sanctions for vegetarianism tend to be mild, but for someone on the fence, they can nevertheless be the difference between making a change and sticking to the status quo. Meat eating is therefore not only supported by other norms but can also be enforced by directly pro-meat-eating social norms.

⁷ While it may seem obvious to many meat eaters that a vegetarian refusing to eat meat is unacceptably violating hospitality norms and being an inconvenience, considering alternative scenarios is illuminating. For example, consider that a host, instead of serving a beef stew, served a dog-meat stew. In this case, it would seem perfectly acceptable for a desire to not eat dog to trump norms of hospitality. Thus the expected order of precedence among norms is itself normatively governed.

The meat eating regularity cannot, therefore, be straightforwardly categorized. For some, eating meat is a personal norm that will be followed even in a meat-hostile environment. Such people may strongly value the flavour of meat and would go out of their way to eat meat even if it were difficult to source. Others eat meat as a custom – it is a straightforward way of meeting nutritional needs that just happens to coincide with what others do. However, many eat meat because they follow a descriptive norm – given that most people in their society eat meat, it makes sense for them to do the same. Meat eating can also be a social norm, one with which most people cheerfully comply (i.e. the norm is internalized and endorsed), but which can force compliance out of reluctant but norm-sensitive individuals. Insofar as normative expectations are a factor in a non-trivial number of people’s meat-eating decisions, a norm-based analysis can aid both in developing an understanding of the current state of affairs and in pointing toward strategies for reducing meat consumption. Although not everyone’s dietary decisions are norm-sensitive, many people’s are (to varying degrees), and people at the margins may change their behaviour in response to relatively small shifts in normative and/or empirical expectations.

4.4 CHANGING NORMS OF MEAT

If the above analysis is correct, then for a significant proportion of the population, meat eating is sensitive to empirical and/or normative expectations. This reinforces the themes of Chapters 2 and 3 – meat eating is a practice sensitive to others’ choices and hence argumentation alone is unlikely change behaviour. Thus if one is interested in transforming meat-eaters into vegetarians, these norm-sensitive social factors must also be addressed. For vegetarianism to become widespread, it would be necessary to undermine current pro-meat norms. However, since meat eaters have independent as well as norm-supported reasons to eat meat, simply undermining pro-meat norms will be insufficient – establishing anti-meat norms will also be required. This section deals with the former; the latter will be addressed in the next section.

4.4.1 *Unravelling pro-meat norms*

There are two variables determining whether someone will be willing to stray from an established norm: their sensitivity to the norm in question, and their attitude to risk with respect to that norm. Norm sensitivity has to do with the number and strength of independent reasons a person has to support a given practice. Someone with high norm sensitivity has internalized the norm and will endorse many reasons to support it (Bicchieri 2017, p. 165). To such a person, some-

one who departs from the norm acts not just imprudently, but *wrongly* (Brennan *et al.* 2013, p. 198). Most meat eaters will put forward reasons like “Meat tastes good” and “It’s easier to eat meat”, but someone highly sensitive to meat-eating norms might also say things like “Animals exist to be eaten”, “It’s not manly to be vegetarian”, or “Vegetarians are trying to destroy the livelihoods of good, hard-working farmers”. For such people, eating meat forms a part of their identity. They will therefore be resistant to information counter to meat eating (Onwezen and Weele 2016; Ruby and Heine 2011). Conversely, someone with low norm sensitivity will have no reasons to follow the norm except reasons tied to the existence of the norm – a fear of disapproval or sanctions.

However, norm sensitivity is not fixed. For example, gaining information about the animal suffering caused by meat production and the availability of alternatives can reduce someone’s norm sensitivity, making them more willing to stop eating meat despite prevalent norms. A reduction in sensitivity to meat norms is therefore an important factor in someone stopping meat consumption, one in which information campaigns can play a role. Nevertheless, this alone is not enough. Whether new information will motivate someone sufficiently for behavioural change will depend on the original level of sensitivity (including strength of norm endorsement), the magnitude of the shift produced by the new information, and the perceived risks of acting counter to the norm.

Risk sensitivity and risk perception together form the other key factor determining if someone will be prepared to abandon an existing norm. Risk sensitivity is an individual’s openness to risk-taking and is a stable personality trait (Bicchieri 2017, p. 174). Some people are naturally more risk-seeking, others are inherently risk-averse. Conversely, risk perception is an individual’s evaluation of the risks of violating a specific norm, and this varies from norm to norm and situation to situation (Bicchieri 2017, p. 174). The more closely the general population adheres to a norm and the stronger the likely sanctions for non-compliance, the greater the perceived (and actual) risks associated with non-compliance (Bicchieri 2017, p. 174). Individual risk perception deviates from the mean both because of errors in risk estimation and because some individuals are more or less vulnerable to sanctions than the average person. These two aspects of risk combine to create a risk attitude for a given person towards a specific norm, which will vary across the population. Thus some people – those who have a favourable risk attitude – will be willing to act earlier than others, if they have reason to.

Attitudes to risk and perceptions of risk affect willingness to transgress a norm. While risk sensitivity is relatively immutable, changes in social conditions can alter perceived risk, and unless perceived risk is addressed, the most sophisticated education campaign

cannot hope to gain widespread success. If the risks of violating a norm are seen as being too high, people will, at most, become persuaded that they ought to act while nevertheless remaining unwilling to do so. Since the perceived normative expectations of others are a proxy for the likelihood of sanctions, normative expectations are an important component of risk perception (Bicchieri 2017, p. 174-175) – stronger normative expectations lead to a higher perceived risk of transgression. Thus the strength of adherence to a norm influences both the perceived and actual risks of transgressing that norm.

Moreover, as I explored in Section 4.2, people tend to care what others think of them even without any perceived risk of sanctions. The risks of violating social norms of meat eating are usually quite low: disapproval and some mockery are likely to be the only social consequences. Nonetheless, these social effects, in combination with the consequences of violating the descriptive norm – the difficulty of finding food one can eat – can be sufficient to counteract weak-to-moderate motivation to change. Fear of others' ill opinion can be sufficient to keep more risk-averse people complying, and a small but vocal minority of meat eaters clearly think badly of vegetarians (Minson and Monin 2012). A candidate for early change must therefore be someone to whom these concerns do not apply, or someone who has less reason than most to care. Hence people who are relatively risk-insensitive with respect to a particular norm are candidates to be early abandoners, if they have reason to – that is, if they also have low norm sensitivity and reasons to transgress.

Bicchieri (2017, p. 163) calls people with low norm sensitivity and low risk perception *trendsetters*. Trendsetters are willing to abandon a norm early and, in doing so, can pave the way for others. In visibly breaking a norm, trendsetters reduce the proportion of people following it and so decrease, at least slightly, the perceived and actual risks others face for violating that norm (Bicchieri 2017, p. 184). If the trendsetter is the right sort of person – someone who has influence either within their social circle or in media – others may also imitate them (Bicchieri 2017, p. 169-170, 194-204). However, whether trendsetters will be able to set change in motion or will simply remain on the fringe depends on just how much their risk perception and norm sensitivity differs from that of others. If their actions reduce perceived risk enough that others follow suit, trendsetters can trigger change. However, if they are not able to reduce the risk of action below the necessary threshold for the next most likely candidate to be confident in changing, then the trendsetters will remain a fringe (Bicchieri 2017, p. 184). With each trendsetter who acts, the bar for the next actor is lowered, and if enough can be persuaded to act, a *tipping point* can be reached, precipitating widespread change.

Because of the role of trendsetters, norms, when they change, do not change linearly. Rather, change starts with a trickle which may

or may not turn into a flood. For any established norm, norm sensitivity and risk perception are jointly sufficient to motivate compliance in most people. However, there will be a few at the margins who could be induced to change their behaviour through a relatively small shift in either their norm sensitivity or their risk perception. This can be seen in Gerry Mackie's (1996) analysis of the decline of foot-binding in China. Contrary to the slow petering out of the process which might have been expected, once change began, the new norm took over within only a few years. There was a generation in which, across most families, the older daughters had their feet bound while younger daughters escaped the practice altogether. The practice of foot-binding was not intrinsically valued by the majority of practitioners, but was rather seen as a necessary evil in order to secure a good marriage for daughters (Mackie 1996). Once a critical mass was reached, the risk to marriage prospects from not binding disappeared, and so the norm could safely be abandoned. Similarly, Anthony Appiah (2010) documents how norms of slavery in the British empire, foot-binding in China, and duelling in England were overturned such that these once-approved practices quickly became deeply moralized. These examples reveal that long-stable norms can be rapidly overturned once momentum is gained; norms are fragile and break readily once non-compliance reaches a critical mass. Meat-eating norms might also change rapidly, if sufficient momentum can be gained.

Not only does each person who either becomes a vegetarian or openly admits that they ought to reduce the perceived risk of vegetarianism, each person who endorses the vegetarian conclusion by their actions also increases the psychic cost of meat eating for remaining meat eaters. The smaller the percentage of people who appear to endorse a particular activity as ethical (or ethically permissible), the harder it is to maintain to oneself that it is (Rabin 1994, p. 179). Fewer people endorsing meat eating would make it harder for remaining meat eaters to continue to maintain the sorts of justifications discussed in Chapter 2. Hence, a sufficient number of people abandoning meat eating could trigger a cascade of change in both behaviour and belief.

Of course, as my discussion of meat eating as a custom revealed, people have independent reasons to want to eat meat, most prominently enjoyment of its taste. This makes meat eating importantly different to a practice like foot-binding, which mothers often struggled with, as is revealed by the Chinese proverb "A mother can't love both her daughter and her daughter's feet at the same time" (Mackie 1996, p. 100). However, the practice of slavery is perhaps more analogous – like meat eating, slavery directly benefited practitioners, despite its moral and human costs. There is therefore precedent for enjoyed practices also being rapidly abandoned and replaced with norms that sanction that practice, as will be examined below. Sim-

ilarly, for widespread vegetarianism to take hold, overturning pro-meat norms would likely prove insufficient: existing norms would need to be replaced by anti-meat norms.

4.4.2 *Establishing anti-meat norms*

In [Chapter 3](#) I argued that both an affective response and a normative framework prohibiting an act are needed to generate a moralized judgement against it. I noted that many aspects of the raising and killing of animals for food have the potential to generate a negative affective response but that the current normative framework encourages desensitization rather than treating this response as morally relevant. I also observed that although it is possible for some individuals to moralize a position based on a personal normative framework, moralization is most likely to succeed where the normative framework has widespread endorsement. I now turn to an exploration of how social norms could come to weigh against meat eating.

Norms currently support meat eating and disincentivise vegetarianism. However, norms can shift; mildly pro-meat norms could be replaced by anti-meat norms. A change in expectations about descriptive norms can sometimes be sufficient to motivate behavioural changes, though this is unlikely to be enough to change behaviour surrounding meat. For example, a series of studies by Cialdini and colleagues (as described in [Bicchieri \(2006, p. 63-70\)](#)) demonstrate that where empirical evidence suggests that people do not litter (i.e. where no litter is present), littering is rare (11-18%). However, where the environment is already dirty, rates of littering increase significantly (32-40%), and they increase even more dramatically where littering behaviour is directly observed (54%). These studies reveal that while some people consistently litter and some never do, many are sensitive to what they expect others to do and adapt their behaviour accordingly even where they have no reason to expect to be seen. Similarly, changes in empirical expectations alone could encourage some people to reduce or eliminate meat consumption.

Social norms can even more strongly reduce or entirely stop participation in activities that people enjoy or otherwise desire, especially where those activities occur in public. For example, in recent decades, smoking has gone from something acceptable, encouraged, and widely enjoyed, to a frowned-upon activity. Its disapproval (at least by peers) dissuades many from taking it up, and can motivate others to stop – a non-smoking social norm has taken hold, albeit imperfectly, in our society. Hence disesteem for violating a social norm and esteem for complying with it can change the cost-benefit ratio of the norm-governed activity, shifting the point at which an activity is rational and thus leading to norm compliance. Esteem considerations can also, over time, lead to norm internalization and moralization. Ini-

tially, anti-meat norms would likely produce mostly grudging conformity. Many might follow the norm in public but violate it in private, and some would flout the norm publicly. However, given time, meat eating could become moralized for most people, and could even become unthinkable. If, as with smoking or littering, people came to expect both that most others would not eat meat and that others would judge them negatively for doing so, then anti-meat norms could take hold.

We eat both in public and in private, and norms have more potential to impact public than private eating. Much food consumption is done privately, either only with family or entirely alone. What is done where nobody can see cannot be subject to esteem considerations – private meat consumption, especially where the meat can be obtained in relatively anonymous ways, will continue, at least in the short term. Similarly, in a family context, meat eating is more likely to persist – we will probably know if our family members desire to eat meat, enabling a different normative context to apply within the home.

Yet meals are frequently communal and social affairs where there is plenty of scope for norms to exert a strong influence, both descriptively through influencing what is available, and normatively by influencing how people think their meat eating will be perceived. In such an environment, norms, modulated by esteem considerations, are likely to have a strong influence. For example, 27% of respondents in our survey agreed they would order a vegetarian meal if dining out with a group who all ordered vegetarian meals (Thrasher and Mayne 2018). Many people, expecting to be disesteemed for eating meat, would avoid it in a public dining setting. Furthermore, were it disesteemed, meat would become less available in restaurants, supermarkets, and other sources of food. This would change people's empirical expectations about the comparative ease of a vegetarian compared to a meat-centred diet.

Over time, norms can shape not just what people do in public, but also how they think. Given a sufficient period of anti-meat norms, it is likely that meat would cease to be seen as food – even by people who are currently enthusiastic defenders of meat eating. With time, reluctant conformity where people refrain from a desired activity to avoid disesteem could transform into sincere norm endorsement and the moralization of meat eating. The raising and killing of animals for meat could become as intuitively immoral as the whipping of downed carthorses.

4.5 MORALIZING MEAT

As highlighted in Chapters 2 and 3, both the justifications put forward by ordinary meat eaters and the ways in which conflicted carnivores

respond to meat reveal that in an important sense, meat eating simply is not treated as a moral issue. Even where it is intellectually acknowledged as such (and often it isn't), both meat-eating behaviour and the justifications put forward in its defence treat meat eating as an issue of personal choice rather than a moral issue with implications for others who matter morally. This section explores why even strict vegetarians often fail to fully moralize meat, and how meat could become widely moralized.

4.5.1 *The current moral status of meat*

The thinkability of an act comes largely from how prevalent it is in one's society; a rational assessment of its badness has little impact. If something just is not done, it can become unthinkable for norm-sensitive people; if it is done, it tends to be thinkable even if one strongly disapproves. As noted in [Chapter 3](#), conflicted carnivores might think that eating meat is objectively worse than petty theft, yet feel a strong aversion to the latter while having no compunctions about the former. Meat eating is highly thinkable.

Yet, as Jean Hampton (2006, p. 109, emphasis original) observes, most people do not think that "constraining ourselves so that we do not rape or murder or steal imposes a *cost* upon us". Not all possible costs enter into ordinary cost-benefit calculations. However, for meat eaters and many people avoiding meat, the costs of giving up meat do factor into calculations. Eating meat is thinkable in a way that murder is not; the forgone opportunities tend to remain salient. Conversely, many vegetarians and vegans re-conceptualize meat so that avoiding it is not seen as a cost – meat is instead seen as something to which they were never entitled. For such people, meat is taken off the table as an option, just as murder or rape are not options, making long-term compliance easier. However, insofar as meat eating is common, it remains thinkable for all but those with a strongly internalized personal norm against it, even among those who disagree that it is morally permissible. If, however, due to a change in norms, meat eating became rare and subject to disapproval, it could be more broadly moved outside the realm of the thinkable.

Currently, even most vegetarians and vegans with strongly internalized personal norms against meat eating and a strong disgust response to meat have not fully moralized meat. Moral rules are generally understood as being universalizable – if something is *morally* wrong for someone to do, then it is wrong for *anyone* to do. Even though extenuating circumstances may make a 'wrong' act all-things-considered acceptable such that the actor is not blameworthy, the act itself nevertheless remains wrong (Nichols 2004, p. 26). Yet the majority of moral vegetarians do not strongly universalize their judgement against meat. Rather they accept that others will eat meat.

Genuinely moralizing a socially sanctioned practice is difficult. On P.F. Strawson's (1962) account of indignation, making a moral judgement requires that one consider it appropriate that others should resent one for failing to act accordingly, and that one feel indignation on observing others failing to comply. Yet it is common for vegetarians to adopt a form of relativism regarding meat – to say that while they think meat eating wrong, it is acceptable for others to eat meat. Furthermore, even those who maintain that meat eating is wrong for all people nevertheless tend to refrain from judging others even if, in some sense, they disapprove. Thus few, if any, vegetarians are making a Strawsonian moral judgement against meat. The strong feelings required by Strawson's conception of moral judgement rely on the disapproved of behaviour not being uncommon – resenting everybody is taxing, just as disesteeming everyone is taxing. Even if meat eating is unthinkable on a personal level, other people's meat eating remains thinkable. Maintaining a strong moral position against a common behaviour is unsustainable.

This disconnect between personal morality and universalizable morality exists in the moral reasoning of both adults and children. There is strong evidence that children intuitively pick up moral rules, knowing that moral rules are universalizable, whereas conventional rules are not (Nichols 2004, p. 6). Yet vegetarian children also fail to universalize their judgements about meat eating.

A study of the moral reasoning of children under the age of ten, comprising of independent vegetarians (children of meat-eating parents who independently decided to become vegetarian), family vegetarians (vegetarian children of vegetarian parents), and non-vegetarians (meat-eating children of meat-eating parents) found that vegetarian children do not properly moralize their vegetarianism (Hussar and Harris 2010). Hussar and Harris (2010) found that all the independently vegetarian children and most of the family vegetarian children cited harm to animals as a motivating factor in their vegetarianism – their reasoning was other-regarding. However, children of all groups judged the meat eating of other people to be morally acceptable. Unsurprisingly, meat-eating children thought their meat eating morally unproblematic. However, both groups of vegetarian children also thought it was morally acceptable to eat meat *as long as the meat eater had not made a commitment to being vegetarian*. Furthermore, all children, including meat-eating children, thought it was morally wrong for a vegetarian to eat meat. Thus rather than focusing on harm to animals, the salient factor for all these children appeared to be whether or not someone was violating their personal moral code. Children would judge people harshly for violating their own code *even when they disagreed with that code*. Thus judgements about the moral permissibility of meat eating seem to rest on individual moral commitments, whereas judgements about

the permissibility of other alleged moral wrongs tend to be treated as deriving from objective moral facts.

We do not ordinarily think that a commitment to not ϕ ing is necessary for ϕ 's wrongness. We think that stealing is wrong, and would not accept a defence along the lines of "I never said I thought stealing was wrong" or "I never committed to not stealing". The same goes for a variety of other wrongs, especially criminal wrongs. Thus it seems that for intuitive wrongness, a prior commitment is necessary for a behaviour against which no norm exist, but where a such a norm exists, the behaviour is treated as *just wrong* irrespective of the commitments of the wrongdoer. Hence as long as meat eating is widespread, it cannot be treated as a proper moral issue; irrespective of personal normative stance, meat will be seen as food.

4.5.2 *Changing the categorization of meat*

'Food' is a normatively governed category. Thus meat could, through norm change, become categorized as not-food. Most potentially nutritious animal flesh is not conceived of as food – what is thinkable as food is shaped by descriptive and social norms. The category of 'edible' animals is much smaller than the category of animals which can be (and in other cultures are) eaten. As Melanie Joy (2010, p. 11-13) argues, our perceptions of meat are far more important than any facts about it. This is why, for most Westerners, a dish believed to be beef can be enjoyed even if it is actually dog, and conversely, why disgust lingers if one has been made to believe that a dish contains dog-meat even if it does not. Thus although most meat eaters think it obvious that the meats they eat are 'food', this categorization is much more malleable than it initially seems.

In any society, there is a small and limited number of animals that are generally recognized as 'edible'. However, this categorization is not based on anything intrinsic to the animals. Which animals are in the 'edible' category and which are not varies between cultures. Nevertheless, strong taboos against eating animals outside the 'edible' category are a constant. Dogs, horses, whales, and pigs are all animals whose meat is enjoyed in some cultures while their meat is regarded with revulsion in others. This cultural variation reveals that the different ways we think (or do not think) about animals affects the way we see their meat. As Hal Herzog (2010) observes, we only eat animals we feel ambivalent towards – we eat neither those we love (pets or particularly cute animals) nor those we hate (pests or particularly ugly animals), with different cultures drawing category lines in different places. While some adventurous eaters take pleasure in sampling flesh from a variety of sources, most people feel uncomfortable eating anything outside a narrow band. For most people, the flesh of an animal that is a 'meat animal' is seen as simply meat; the

flesh of other animals cannot be separated from the animal that was (Joy 2010, p. 11-13), and thus cannot be seen as 'meat'. Norms thus inform which animals can be eaten, and which cannot even be seen as food.

Given that what is 'food' is a norm-governed category, compared to what is 'in theory edible', a shift in norms could take meat 'off the table'. This may at first seem implausible, but many practices that were once both normatively encouraged and enjoyed are now not just seen as wrong, they are effectively unthinkable. Practices ranging from ritualistic cat burning to bear baiting to public corporal punishment were not just cruel practices employed in the service of goals such as deterrence of criminals, they were also enjoyed as participatory spectacles when they served no greater purpose (Pinker 2011, p. 144-148). Yet as Nichols (2004, p. 143-144) documents, anti-harm norms have consistently been expanding over the centuries to cover ever more categories of people, as well as many animals. Violence is becoming less and less accepted. When a new group is added to the list of those against whom violence is unacceptable, or a new practice becomes understood as harmful, norm change tends to happen abruptly, on the scale of years rather than generations. Support for practices newly categorized as harmful has therefore dropped abruptly and turned to revulsion once a critical mass of society accepted that re-categorization. This change can be so extreme that people who enjoyed participation in the above-described practices in their youth frequently came to later regard them not just as wrong, but as incomprehensible (Pinker 2011, ch. 7). Norms can reverse to the point where the original normatively encouraged behaviour is deeply moralized.

If a critical mass of people were to abandon meat eating, vegetarianism could become 'normal' – as it is in parts of India, if for different reasons to those which drive most Western vegetarians.⁸ As long as meat is common, it is thinkable and cannot be properly moralized. If, however, it were to cease to be normal, then given the aversion people have to seeing suffering, it is probable that, like cat burning, pig eating would also become unthinkable.

Thus far, I have argued that given the right normative environment, meat eating could become a deeply moralized behaviour. Of course, transitioning to such an environment will be difficult and may seem implausible, especially as vegetarian movements have long existed but have gained little traction. I do not attempt here to provide a road map to norm change – that is a task for psychology or sociology. Nevertheless, it is worth briefly demonstrating that widespread change for the sake of animals can occur even where this inconveniences the people making the change.

⁸ Around 40% of all Indian households are lacto-vegetarian (Delgado, Narrod, and Tiongco 2003, §2.3), and among certain subsets of the population, particularly adherents of Jainism, as well as many geographic groups of Brahmins, vegetarianism is both normatively and empirically expected.

One example, documented by Martin Balluch (2006), is an Austrian ban on eggs from caged hens, implemented with popular support. Animal activists were able to leverage widespread disapproval of the conditions under which battery hens lived (opposed by 86% of the population) to gain general support for a battery hen ban. Importantly, this ban was supported even by people who bought cage eggs – there was a willingness to sacrifice a certain level of personal convenience given assurance that this sacrifice would be effective in securing better welfare outcomes for Austrian chickens. In this case, an assurance of collective action, achieved through legal change, improved conditions for animals.

Of course, a collective agreement to ban meat production is extremely implausible in the current environment in any Western nation. Nevertheless, it is possible that, given time and the right impetus, meat eating may become restricted to subcultures or eliminated altogether. My analysis has revealed that if one is interested in changing meat-eating behaviour, multiple factors need to be addressed. Sensitivity to existing norms of meat eating need to be reduced and positive reasons to avoid meat adopted. These changes can be aided by information campaigns and even argumentation. However, such an approach will be ineffective unless complemented by efforts to address both the practical and social barriers to avoiding meat. Practical barriers can be addressed by increased availability of vegetarian options. While this change will depend in part on demand, other restrictions on the prevalence and desirability of vegetarian food can be addressed by improved alternatives to farmed meat, including high-quality plant-based mock meats or even cultured meat. Similarly, while many of the social factors disincentivising meat avoidance are dependent on the number of vegetarians in a society, other social pressures can be alleviated by changing attitudes towards vegetarians and vegetarianism without necessarily attempting to change attitudes to meat. These are tasks for activists and entrepreneurs, but they can nevertheless be approached in ways informed by this philosophical analysis.

4.6 CONCLUSION

Meat eating is common in most societies, and where it is not common, it is generally an aspiration. It is normal – eating meat is what most people do – but more than this, it is a norm. Both the empirical and normative expectations people hold contribute to their meat-eating behaviour. Thus norms of meat eating make vegetarianism challenging at multiple levels. The normalcy of meat makes it easy for people to justify eating the meat of intensively farmed animals, even where ignoring the suffering involved goes against their broader moral beliefs (Chapter 2). The normalcy of meat also makes it difficult to mo-

tivate behavioural change even where one thinks one should ([Chapter 3](#)), and even if one is motivated to change, the normalcy of meat can make it difficult to sustain that change.

While most people value being moral (or being perceived to be so), this incentive is tempered by self-interested considerations. In the case of meat eating, such considerations are most prominently convenience and taste. These considerations need to be traded against the value of acting morally, and the way one expects others to perceive one's choices plays an important role in how that trade-off is understood. Given that most people eat meat, one can be sufficiently moral according to societal standards while continuing to eat meat, making motivating vegetarianism difficult. Hence norms surrounding meat eating shape the potential for people to see meat eating as a genuine moral issue, and thus the likelihood they will be motivated to act.

According to Cristina Bicchieri's norm framework, regularities in behaviour can be motivated in different ways. Customs are followed because a particular behaviour is the best way for individuals to meet their needs independently of what others are doing. Descriptive norms receive conditional adherence based on what others are doing; a descriptive norm makes conformity the best way of reaching one's goals insofar as one expects others to also conform. In contrast, social norms rely on both normative and empirical expectations. It is rational to conform to a social norm, even if one dislikes it, if one both expects that others will conform (empirical expectations) and also expects that others expect that one should conform (second order normative expectations).

While at first glance meat eating may seem to be merely a custom, it is not – both empirical and normative expectations play a non-trivial role in sustaining the practice. First, the behaviour of others shapes expectations about the relative ease or difficulty of obtaining meat-containing and vegetarian meals. Second, normative expectations shape which animals one can and cannot eat. The flesh of some animals (in mainstream Western society, mostly fish, cows, chickens, and pigs) is seen as 'meat', whereas the flesh of others (for example, horses, dogs, and whales) is seen as inedible. Third, many vegetarians and vegetarian-sympathetic people report social pressure to eat meat – normative expectations make it difficult to not eat meat, particularly when one is eating amongst others. While meat eating is merely a custom for some, it is a descriptive norm for others, and a social norm for many.

Thus attempts to motivate vegetarianism that do not address the relevant normative and empirical expectations are highly unlikely to succeed. This is why merely refuting the justifications for meat eating discussed in [Chapter 2](#) tends to be insufficient for changing people's behaviour. For practical persuasion, it is necessary to address both what one expects others to do, and what one expects others to expect

one to do. The people most likely to be receptive to acting on a pro-vegetarian message – potential trendsetters – will be those who have both low norm sensitivity (i.e. people who do not have strong reasons to eat meat and who are receptive to anti-meat messages) and who also perceive the risks of vegetarianism to be low (i.e. people who expect relatively little social fallout and who think that they will be able to get decent food as vegetarians). This may be why students are particularly receptive to pro-vegetarian messages – many have moved out of their parents' home and are no longer under pressure to fit in with their family's meals, but also have not yet developed relationships or become parents, both of which make change more difficult and risky. The key point, however, is that approaches that focus on education (i.e. reducing norm sensitivity) while neglecting risk perceptions are unlikely to gain much traction.

Because categories like 'food' are normatively constructed, if sufficient people adopt vegetarian norms, then a cascade could be triggered, inducing vegetarianism even in people who are currently reluctant to cease eating meat. Although vegetarianism has existed as a fringe movement throughout most of human history, rarely reaching this tipping point, if such a tipping point is reached, it is likely that aversion to harm will result in the moralization of eating the flesh of once-living animals.

CONCLUSION

The way the majority of meat eaters think about meat is incoherent. Most meat eaters' justifications are inconsistent with their broader moral beliefs – the justifications do not justify, but rather seek to excuse a desired activity. Many meat eaters actively avoid contemplating what is involved in producing the meat they eat. Others freely acknowledge that it is immoral for them to eat meat, but do so anyway, without experiencing any significant guilt. The suffering of animals for meat production is arguably a serious moral issue, but, as with comparable issues such as extreme poverty and climate change, it is rarely seen as an issue that demands personal action. This thesis has explored the moral psychological reasons underpinning both the widespread failure to recognize meat eating as a moral issue, and the continued consumption of meat even by most people who recognize it as a moral issue. I have argued that the inconsistencies in how people reason about meat are deeply tied to meat's normalcy – both in terms of its being common and its being supported and enforced by descriptive and social norms.

As someone who hopes to reduce human-inflicted animal suffering, instead of following the well-trodden path of presenting ethical arguments for an obligation to stop eating meat, I have pursued what I hope is a more productive path: seeking to understand why those arguments tend to be ineffective. I have therefore explored why it is that some issues are treated as being deeply moral such that 'doing the wrong thing' is effectively unthinkable, why other issues, if they are even recognized as being wrong, are treated as violable, and why meat eating tends to fall into the latter category. In doing so, I have shown that argumentative persuasion targeting individuals is unlikely to be effective; social factors rather than individual beliefs about morality shape not only willingness to act on an issue, but also ability to recognize an issue as a moral one demanding action. The understanding of how people reason about moral issues gained through my analysis should be illuminating both to people interested specifically in the issue of meat eating and those interested in moral reasoning more broadly.

In [Chapter 2](#), I argued that meat eating seems easy to justify not because the available justifications succeed, but because meat eating is common. I examined the lay justifications for meat eating and argued that they were all inadequate for any meat eater who endorses the *meat trilemma*:

1. *ceteris paribus*, causing animals to suffer is morally bad because of what it does to the animal – animals have moral status;
2. *ceteris paribus*, one should try to be a moral person and not act in ways that have morally bad results; and
3. eating meat, including the meat of animals who were raised on factory farms or can otherwise be reasonably inferred to have led bad lives, is morally permissible.

Strong necessity requires denial of readily available information about the potential for a healthy vegetarian lifestyle: it relies on a manifestly false belief, at least in the general case. All the other lay justifications – weak necessity, normalcy, naturalness, niceness, and no effect – require treating meat eating differently to issues that are widely accepted as being moral ones. They all require accepting reasoning that would have been rejected when considering a more clearly moral issue such as theft or murder. In other words, they assume meat's moral permissibility in order to establish it. None can resolve the meat paradox.

I then examined the more sophisticated philosophical arguments and found that they fare no better. Most sophisticated defences of meat eating explicitly rule out defending the most commonly consumed type of meat – meat from intensively raised animals who can reasonably be inferred to have experienced non-trivial levels of suffering. Thus these defences are not available to any meat eater who thinks animals matter morally and who at least sometimes eats the meat of such animals. Another defence could only justify eating meat where it could reasonably be expected to provide an exceptional dining experience. Those philosophical arguments that either directly or indirectly defend eating the meat of animals irrespective of the animals' quality of life required that *all* animals be treated as having no moral status in their own right. Given, however, that most people think that mistreating animals can at least sometimes be wrong because of what it does to the animal, I argued that this is not a bullet most meat eaters would be prepared to bite. Thus the best available philosophical arguments also fail to justify meat eating for ordinary meat eaters.

Despite these inadequacies, however, most meat eaters do not experience guilt regarding their meat eating. I argued that justifications like those I examined, along with strategic avoidance of unwanted information, play a role in managing cognitive dissonance. Because such justifications are widely accepted, they can be used to maintain the belief that eating animals is morally acceptable, even though this conflicts with other beliefs. This causes some cognitive load. However, due to meat eating's prevalence, these justifications are rarely challenged, and so little cognitive load is incurred. This mechanism of dissonance management is further reinforced by strategic ignorance

– deliberately avoiding information that might render the plight of farmed animals salient, or make one’s justifications for eating meat seem flimsy. Thus good-conscience meat eating is, I argued, dependent on its being widespread.

In [Chapter 3](#) I turned to exploring the puzzle of the *conflicted carnivore*. Conflicted carnivores, by definition, agree that morally, they ought not eat meat. They do so nonetheless, and experience little ensuing guilt. I argued that irrespective of whether the conflicted carnivore’s judgement against meat is perceived as sincere, there is a motivational gap to be accounted for: why is it that the conflicted carnivore makes sincere or motivating moral judgements in some domains but does not include meat eating in this class of judgements? I argued that the conflicted carnivore’s judgement about meat’s impermissibility is not being moralized – the intellectual judgement that eating meat is wrong is not converted into a *feeling* of wrongness. I argued that because of our socially endorsed normative structure, affectively upsetting stimuli connected to meat are treated as something to be avoided rather than addressed; negative affective responses to meat are not treated as carrying moral information. Hence meat is not being moralized.

Following Shaun Nichols (2004), I argued that a moralized judgement requires two components: a negative affective response and a normative framework indicating that the thing under judgement is wrong. While it is possible to moralize something based on an individual moral norm, socially endorsed rules tend to play this role, and our society-wide rules do not say that eating animals is wrong – indeed, they say the opposite. Thus an intellectual judgement that meat eating is wrong tends to be supported neither by a strong affective response nor by a society-level normative rule. Instead, because the rule is absent, the affective response, where present, tends to be suppressed. I argued that this enables meat eating to seem morally acceptable despite intellectual judgements to the contrary. Thus for conflicted carnivores, too, the fact that meat eating is both widespread and widely accepted makes it easy to ignore one’s intellectual moral judgement against meat eating.

Hence in [Chapter 4](#), I turned to an analysis of the behavioural regularity of meat eating. I sought to understand how social dynamics shape the way people currently think about meat and what could change this. I argued that while meat eating might seem to be simply the collective expression of individual preferences (what Cristina Bicchieri (2017) calls a custom), it is, in most cases, a behaviour heavily influenced by social factors. People eat the meat of only a small fraction of potentially edible species, and generally view the meat of other species not only with suspicion, but with revulsion. What is edible and what is not are normatively governed. Furthermore, the fact that meat eating is widespread whereas vegetarianism is relatively

uncommon makes it easy to obtain meals containing meat but difficult to source tasty and nutritious vegetarian meals, and even harder to find vegan ones. The expectation that others will eat meat therefore leads to the empirical expectation that one should also eat meat if one wants eating to be straightforward. Thus meat eating has aspects of what Bicchieri calls a descriptive norm. Furthermore, a normative aspect often attaches to meat eating. People can experience implicit or explicit social pressure to conform to what most people eat, especially in social situations where eating differently can inconvenience one's host. Thus many people experience meat eating as a social norm. This norm is usually internalized – a norm with which one agrees and wants to comply – but it can also exert unwanted power on a weak-willed conflicted carnivore.

Meat's status as a norm is a double-edged sword – it currently helps to entrench meat eating, but could work against it. Given, however, that eating meat also satisfies a personal preference, simply overturning pro-meat norms would be insufficient – anti-meat norms would also need to be established. For the abandonment of a 'bad' norm, it is ineffective to simply persuade individuals that they should abandon it, as should be obvious from the number of conflicted carnivores who have been persuaded but do not act. Changing people's behaviour therefore requires acting on both norm sensitivity (the range and strength of reasons people have to support the norm) and risk (the perceived risk of deviating from the norm). The specifics of how to achieve this are a task for psychologists and social scientists, not philosophers. Nevertheless, I have highlighted how such research might be productively directed.

In arguing for the role played by norms, I have not given a merely sociological explanation of why people eat meat. Rather, in showing how people reason about meat eating specifically, and moral issues more generally, I have uncovered important implications for how philosophers ask questions about moral behaviour in general. The way we look at moral 'failures' should not ask questions only about individual factors such as rationality, sincerity, affect, or will. The society in which a person is embedded must also be examined. Focusing on the failures of abstracted individuals misses essential insights into what motivates moral behaviour. I have shown that individual moral decision-making is deeply embedded in a wider social environment. How other people behave with respect to an issue affects both individual willingness to even entertain it as an issue of moral import, and also one's ability to convince oneself to act on an issue that has intellectually been recognized as morally significant. Thus the way we as philosophers ask questions about moral reasoning needs to be reassessed – it needs to take social factors into account while continuing to recognize individuals as the locus of moral decision-making.

My analysis also has implications for the practice of ethical argument. If people are not engaging with arguments and are not motivated to change their behaviour even where they are intellectually persuaded, arguments such as those presented by ethicists are unlikely to change minds. Thus although arguments in ethics may be valuable as philosophical exercises, engaging in ethical argument is not a fruitful pastime if the goal is changing minds and behaviour. Thus philosophers who are engaging in ethical argument *with the hope of changing minds* would do better to apply philosophical scrutiny to broader issues surrounding the culture of the issue in which they are interested.

Though I have explored the discrepancies in how people think about meat eating, such discrepancies are not unique to this issue. It is likely that people rationalize their inaction on other issues such as failing to act on climate change or extreme poverty in ways similar to those explored in [Chapter 2](#), and they do so for similar reasons. People also fail to act on their intellectual moral judgements about these other issues, probably for reasons like those explored in [Chapter 3](#). Similarly, it is likely that the same or similar social factors to those explored in [Chapter 4](#) influence attitudes and behaviour regarding these other issues. Thus while my analysis focused on meat eating, someone interested in understanding the moral psychology of other comparable issues could adapt my approach, likely finding many similarities (as well as some interesting differences) between how people think about meat eating and how they think about these similar moral issues.

Individuals make judgements and choices, but they do so within a social environment. While personal moral beliefs can play a role, they are frequently overwhelmed by factors in one's social environment. My analysis of the moral psychology of meat eating reveals that someone not being motivated to stop eating meat is not just a failure of reason, affect, or motivation. Rather, it is at least a partial product of one's social environment, which can make it rational to conform to what others are doing irrespective of one's personal beliefs about what morality requires. Furthermore, one's social environment makes it difficult to achieve or maintain moral clarity about an issue that is not widely treated as a moral one. It is likely that one's social environment plays a similar role in other issues I have not explicitly discussed. Ignoring one's role in climate change may be similarly enabled by the fact that few people act to reduce their emissions. Similarly, the fact that most people do not give significantly to charities and buy cheap products without consideration for how they were manufactured likely reassures people that their own behaviour in this domain is morally acceptable. While there are subtle differences between these cases, all are issues which are largely ignored in mainstream Western society, meaning social pressures to behave 'well'

are largely absent. Thus understanding the causes and effects of these social pressures is essential to understanding how people respond to moral issues.

BIBLIOGRAPHY

- Adams, C (2000). *The sexual politics of meat: a feminist-vegetarian critical theory*. 10th anniv. Continuum, New York.
- Allen, MW and Baines, S (2002). "Manipulating the symbolic meaning of meat to encourage greater acceptance of fruits and vegetables and less proclivity for red and white meat." *Appetite* vol. 38, no. 2, pp. 118–130. DOI: [10.1006/appe.2001.0474](https://doi.org/10.1006/appe.2001.0474).
- American Dietetic Association (2016). "Position of the American Dietetic Association: vegetarian diets." *Journal of the American Dietetic Association* vol. 116, no. 12, pp. 1970–1980. DOI: [10.1016/j.jand.2016.09.025](https://doi.org/10.1016/j.jand.2016.09.025).
- Appiah, A (2010). *The honor code: how moral revolutions happen*. W.W. Norton, New York.
- Bagemihl, B (1999). *Biological exuberance: animal homosexuality and natural diversity*. Profile Books, London.
- Baier, K (1965). *The moral point of view: a rational basis of ethics*. Cornell University Press, Ithaca, New York.
- Balluch, M (2006). "How Austria achieved a historic breakthrough for animals." In: *In defense of animals: the second wave*. Ed. by P Singer. Malden, MA, Blackwell Publishing, pp. 157–166.
- Bastian, B, Loughnan, S, Haslam, N, and Radke, H (2012). "Don't mind meat? the denial of mind to animals used for human consumption." *Personality and Social Psychology Bulletin* vol. 38, no. 2, pp. 247–256. DOI: [10.1177/0146167211424291](https://doi.org/10.1177/0146167211424291).
- BBC (2008). "China bans dog from Olympic menu." *BBC News*. URL: <http://news.bbc.co.uk/2/hi/asia-pacific/7501768.stm>.
- Belshaw, C (2016). "Meat." In: *The moral complexities of eating meat*. Ed. by B Bramble and B Fischer. New York, Oxford University Press, pp. 9–29.
- Berndsen, M and Van Der Pligt, J (2004). "Ambivalence towards meat." *Appetite* vol. 42, no. 1, pp. 71–78. DOI: [10.1016/S0195-6663\(03\)00119-3](https://doi.org/10.1016/S0195-6663(03)00119-3).
- Bicchieri, C (2006). *The grammar of society: the nature and dynamics of social norms*. Cambridge University Press, New York.
- Bicchieri, C (2017). *Norms in the wild*. Oxford University Press, New York.
- Bicchieri, C and Xiao, E (2009). "Do the right thing: but only if others do so." *Journal of Behavioral Decision Making* vol. 22, pp. 191–208. DOI: [10.1002/bdm.621](https://doi.org/10.1002/bdm.621).
- Bramble, B and Fischer, B, eds. (2016). *The moral complexities of eating meat*. Oxford University Press, New York.

- Brennan, G, Eriksson, L, Goodin, R, and Southwood, N (2013). *Explaining norms*. Oxford University Press, Oxford.
- Brink, D (1989). *Moral realism and the foundations of ethics*. Cambridge University Press, Cambridge.
- Bruckner, D (2016). "Strict vegetarianism is immoral." In: *The moral complexities of eating meat*. Ed. by B Bramble and B Fischer. New York, Oxford University Press, pp. 30–47.
- Budolfson, MB (2016). "Is it wrong to eat meat from factory farms? if so, why?" In: *The moral complexities of eating meat*. Ed. by B Bramble and B Fischer. New York, Oxford University Press, pp. 80–98.
- Carruthers, P (1992). *The animals issue: moral theory in practice*. Cambridge University Press, Cambridge, UK.
- Cialdini, RB, Demaine, LJ, Sagarin, BJ, Barrett, DW, Rhoads, K, and Winter, PL (2006). "Managing social norms for persuasive impact." *Social Influence* vol. 1, no. 1, pp. 3–15. DOI: [10.1080/15534510500181459](https://doi.org/10.1080/15534510500181459).
- Criminal Code Act (1995). c. 2. URL: http://www9.austlii.edu.au/cgi-bin/viewdoc/au/legis/cth/consol{_}act/cca1995115/sch1.html.
- Cuneo, T (2016). "Conscientious omnivorism." In: *Philosophy comes to dinner: arguments about the ethics of eating*. Ed. by A Chignell, T Cuneo, Halteman, and M C. New York, Routledge, pp. 21–38.
- Dana, J (2005). "Conflicts of interest and strategic ignorance of harm." In: *Conflicts of interest: challenges and solutions in business, law, medicine, and public policy*. Ed. by DA Moore, D Cain, G Loewenstein, and M Bazerman. Cambridge, Cambridge, pp. 206–223.
- Dana, J, Weber, RA, and Kuang, JX (2007). "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness." *Economic Theory* vol. 33, no. 1, pp. 67–80. DOI: [10.1007/s00199-006-0153-z](https://doi.org/10.1007/s00199-006-0153-z).
- DeGrazia, D (2009). "Moral vegetarianism from a very broad basis." *Journal of Moral Philosophy* vol. 6, no. 2, pp. 143–165. DOI: [10.1163/174552409X402313](https://doi.org/10.1163/174552409X402313).
- Delgado, CL, Narrod, CA, and Tiongco, MM (2003). *Project on livestock industrialization, trade and social-health-environment impacts in developing countries*. Tech. rep. Food and Agriculture Organization of the United Nations. URL: <http://www.fao.org/WAIRDOCS/LEAD/X6115E/x6115e00.htm{\#}Contents>.
- Descartes, R (1637). "Discourse on method." In: *The philosophical writings of Descartes*. Ed. by J Cottingham, R Stoothoff, and D Murdoch. Trans. by R Stoothoff. Cambridge, 1985, Cambridge University Press, pp. 109–77.
- Dietitians Association of Australia (2016a). *Vegan diets*. URL: <http://daa.asn.au/for-the-public/smart-eating-for-you/nutrition-a-z/vegan-diets/> (visited on 01/03/2016).

- Dietitians Association of Australia (2016b). *Vegetarian diets*. URL: <http://daa.asn.au/for-the-public/smart-eating-for-you/nutrition-a-z/vegetarian-diets/> (visited on 01/03/2016).
- Driver, J (2016). "Individual consumption and moral complicity." In: *The moral complexities of eating meat*. Ed. by B Bramble and B Fischer. New York, Oxford University Press, pp. 67–79.
- Eagly, AH and Chaiken, S (1993). *The psychology of attitudes*. Harcourt Brace Jovanovich College Publishers, Fort Worth.
- Elster, A and Miller, JE (2008). *I still see her haunting eyes: the holocaust & a hidden child named Aaron*. 3rd. BF Press, Peoria, IL.
- Engel, MJ (2000). "The immorality of eating meat." In: *The moral life: an introductory reader in ethics and literature*. Ed. by LP Pojman. New York, Oxford University Press.
- Feinberg, J (1984). *The moral limits of the criminal law, volume 1: harm to others*. Oxford University Press, New York.
- Finnis, J (2015). "Natural law theories." *Stanford Encyclopedia of Philosophy*. URL: <https://plato.stanford.edu/entries/natural-law-theories/>.
- Fischer, B (2018). "Arguments for consuming animal products." In: *The Oxford handbook of food ethics*. Ed. by A Barnhill, M Budolfson, and T Doggett. New York, Oxford University Press. Chap. 11.
- Foer, JS (2009). *Eating animals*. Black Bay Books, New York.
- Gibbard, A (1990). *Wise choices, apt feelings: a theory of normative judgment*. Harvard University Press, Cambridge, Massachusetts.
- Greene, JD (2009). "Dual-process morality and the personal/impersonal distinction: A reply to McGuire, Langdon, Coltheart, and Mackenzie." *Journal of Experimental Social Psychology* vol. 45, no. 3, pp. 581–584. DOI: [10.1016/j.jesp.2009.01.003](https://doi.org/10.1016/j.jesp.2009.01.003).
- Haidt, J (2001). "The emotional dog and its rational tail: a social intuitionist approach to moral judgment." *Psychological Review* vol. 108, no. 4, pp. 814–834. DOI: [10.1037//0033-295X](https://doi.org/10.1037//0033-295X).
- Haidt, J (2006). *The happiness hypothesis: finding modern truth in ancient wisdom*. Basic Books, New York.
- Hamad, R (2013). "Halal." In: *Defiant daughters: 21 women on art, activism, animals, and the sexual politics of meat*. Ed. by K Davis and W Lee. New York, Lantern Books.
- Hampton, J (2006). "Righting wrongs: the goal of retribution." In: *The intrinsic worth of persons: contractarianism in moral and political philosophy*. Ed. by D Farnham. Cambridge, UK, Cambridge University Press, pp. 108–150.
- Herron, M (2015). *The many meanings of 'free range' chicken*. URL: <https://www.choice.com.au/food-and-drink/meat-fish-and-eggs/meat/articles/free-range-chicken> (visited on 05/29/2017).
- Herzog, H (2010). *Some we love, some we hate, some we eat: why it's so hard to think straight about animals*. Harper, New York.

- Hsiao, T (2015). "In defense of eating meat." *Journal of Agricultural and Environmental Ethics* vol. 28, no. 2, pp. 277–291. DOI: [10.1007/s10806-015-9534-2](https://doi.org/10.1007/s10806-015-9534-2).
- Hsiao, T (2017). "Industrial farming is not cruel to animals." *Journal of Agricultural and Environmental Ethics* vol. 30, no. 1, pp. 37–54. DOI: [10.1007/s10806-017-9652-0](https://doi.org/10.1007/s10806-017-9652-0).
- Humane Research Council (2014). *Study of current and former vegetarians and vegans – tables & methodology*. URL: https://faunalytics.org/wp-content/uploads/2015/06/Faunalytics{_}Current-Former-Vegetarians{_}Full-Report.pdf.
- Hume, D (1739). "A treatise of human nature: of the passions." In: *David Hume: moral philosophy*. Ed. by G Sayre-McCord. Indianapolis, 2006, Hackett Publishing, pp. 13–66.
- Hume, D (1740). "A treatise of human nature: of morals." In: *David Hume: moral philosophy*. Ed. by G Sayre-McCord. Indianapolis, 2006., Hackett Publishing, pp. 67–184.
- Hursthouse, R (2006). "Applying virtue ethics to our treatment of the other animals." In: *The practice of virtue: classic and contemporary readings in virtue ethics*. Ed. by J Welchman. Indianapolis, Hackett Publishing, pp. 136–155.
- Hussar, KM and Harris, PL (2010). "Children who choose not to eat meat: a study of early moral decision-making." *Social Development* vol. 19, no. 3, pp. 627–641. DOI: [10.1111/j.1467-9507.2009.00547.x](https://doi.org/10.1111/j.1467-9507.2009.00547.x).
- Joy, M (2010). *Why we love dogs, eat pigs, and wear cows*. Conari Press, San Francisco.
- Kelly, D (2011). *Yuck! the nature and moral significance of disgust*. Vol. 2011. The MIT Press, Cambridge, Mass. DOI: [10.1080/09515089.2012.696329](https://doi.org/10.1080/09515089.2012.696329).
- Kennett, J and Fine, C (2009). "Will the real moral judgment please stand up? the implications of social intuitionist models of cognition for meta-ethics and moral psychology." *Ethical Theory and Moral Practice* vol. 12, no. 1, pp. 77–96.
- Korsgaard, CM (2004). "Fellow creatures: Kantian ethics and our duties to animals." *The Tanner Lecture on Human Values*. URL: <http://nrs.harvard.edu/urn-3:HUL.InstRepos:3198692>.
- Lea, E and Worsley, A (2003). "Benefits and barriers to the consumption of a vegetarian diet in Australia." *Public Health Nutrition* vol. 6, no. 5, pp. 505–11. DOI: [10.1079/PHN2002452](https://doi.org/10.1079/PHN2002452).
- Lewis, D (2002). *Convention*. Blackwell Publishers, Oxford.
- Lipscomb, BJB (2016). "'Eat responsibly': agrarianism and meat." In: *Philosophy comes to dinner: arguments about the ethics of eating*. Ed. by A Chignell, T Cuneo, and MC Halteman. New York, Routledge, pp. 56–72.

- Lomasky, L (2013). "Is it wrong to eat animals?*" *Social Philosophy and Policy* vol. 30, no. 1-2, pp. 177–200. DOI: [10 . 1017 / S0265052513000083](https://doi.org/10.1017/S0265052513000083).
- Loughnan, S, Bastian, B, and Haslam, N (2014). "The psychology of eating animals." *Current Directions in Psychological Science* vol. 23, no. 2, pp. 104–108. DOI: [10.1177/0963721414525781](https://doi.org/10.1177/0963721414525781).
- Mackie, G (1996). "Ending footbinding and infibulation: a convention account." *American Sociological Review* vol. 61, no. 6, p. 999. DOI: [10.2307/2096305](https://doi.org/10.2307/2096305).
- Mackie, JL (1977). *Ethics: inventing right and wrong*. Penguin, New York.
- Mason, J and Singer, P (1990). *Animal factories*. Revised. Three Rivers Press, New York.
- McCurry, J (2017). "South Korea closes biggest dog meat market in run-up to Olympics." *The Guardian*. URL: <https://www.theguardian.com/world/2017/feb/28/south-korea-closes-biggest-dog-meat-market-in-run-up-to-olympics>.
- McDowell, J (1998). *Mind, value, and reality*. Harvard University Press, Cambridge, MA.
- McPherson, T (2016). "Why I am a vegan (and you should be one too)." In: *Philosophy comes to dinner*. Ed. by A Chignell, T Cuneo, and MCH Corporation. London, Routledge, pp. 73–75.
- Mercier, H and Sperber, D (2011). "Why do humans reason? arguments for an argumentative theory." *Behavioral and Brain Sciences* vol. 34, no. 02, pp. 57–74. DOI: [10.1017/S0140525X10000968](https://doi.org/10.1017/S0140525X10000968).
- Minson, J and Monin, B (2012). "Do-gooder derogation: disparaging morally motivated minorities to defuse anticipated reproach." *Social Psychological and Personality Science* vol. 3, pp. 200–207. DOI: [10.1177/1948550611415695](https://doi.org/10.1177/1948550611415695).
- Murphy, SP and Allen, LH (2003). "Animal source foods to improve micronutrient nutrition and human function in developing countries." *Source* vol. 1933, no. 5, pp. 3932–3935. DOI: [0022 - 3166 / 03](https://doi.org/10.1002-3166/03). arXiv: [03 \[0022-3166\]](https://arxiv.org/abs/030022-3166).
- Nichols, S (2004). *Sentimental rules*. Oxford University Press, Oxford.
- Onwezen, MC and Weele, CN van der (2016). "When indifference is ambivalence: strategic ignorance about meat consumption." *Food Quality and Preference* vol. 52, pp. 96–105. DOI: [10.1016/j.foodqual.2016.04.001](https://doi.org/10.1016/j.foodqual.2016.04.001).
- Pettit, P (2002). *Rules, reasons, and norms*. Clarendon Press, Oxford.
- Pettit, P and Brennan, G (2005). *The economy of esteem*.
- Piazza, J, Ruby, MB, Loughnan, S, Luong, M, Kulik, J, Watkins, HM, and Seigerman, M (2015). "Rationalizing meat consumption. The 4Ns." *Appetite* vol. 91, pp. 114–128. DOI: [10.1016/j.appet.2015.04.011](https://doi.org/10.1016/j.appet.2015.04.011).
- Pinker, S (2011). *The better angels of our nature: why violence has declined*. Penguin Group, New York.

- Plous, S (1993). "Psychological mechanisms in the human use of animals." *Journal of Social Issues* vol. 49, no. 1, pp. 11–52.
- Pollan, M (2006). *The omnivore's dilemma: a natural history of four meals*. Penguin Press, New York.
- Pollan, M (2009). *The omnivore's dilemma: the secrets behind what you eat*. Dial Books for Young Readers, New York.
- Rabin, M (1994). "Cognitive dissonance and social change." *Journal of Economic Behavior and Organization* vol. 23, no. 2, pp. 177–194. DOI: [10.1016/0167-2681\(94\)90066-3](https://doi.org/10.1016/0167-2681(94)90066-3).
- Rachels, J (1997). "The moral argument for vegetarianism." In: *Can ethics provide answers? and other essays in moral philosophy*. Lanham, Rowman & Littlefield, pp. 99–107.
- Reese, J (2017). "Survey of US attitudes towards animal farming and animal-free food October 2017." *Sentience Institute*. URL: <https://www.sentienceinstitute.org/animal-farming-attitudes-survey-2017>.
- Regan, T (2004). *The case for animal rights*. 2nd. University of California Press, Berkeley.
- Rosati, CS (2014). "Moral motivation." *The Stanford Encyclopedia of Philosophy*. URL: <http://plato.stanford.edu/archives/spr2014/entries/moral-motivation/>.
- Rothgerber, H (2012). "Real men don't eat (vegetable) quiche: masculinity and the justification of meat consumption." *Psychology of Men & Masculinity* vol. 14, no. 4, pp. 1–13. DOI: [10.1037/a0030379](https://doi.org/10.1037/a0030379).
- Rothgerber, H (2014). "Efforts to overcome vegetarian-induced dissonance among meat eaters." *Appetite* vol. 79, pp. 32–41. DOI: [10.1016/j.appet.2014.04.003](https://doi.org/10.1016/j.appet.2014.04.003).
- Rowlands, M (2002). "Justice for all." In: *Animals like us*. London, Verso. Chap. 3.
- Royzman, EB, Kim, K, and Leeman, RF (2015). "The curious tale of Julie and Mark: unraveling the moral dumbfounding effect." *Judgment and Decision Making* vol. 10, no. 4, pp. 296–313.
- Rozin, P (1999). "The process of moralization." *Psychological Science* vol. 10, no. 3, pp. 218–221.
- Rozin, P, Markwith, M, and Stoess, C (1997). "Moralization and becoming a vegetarian: the transformation of preferences into values and the recruitment of disgust." *Psychological Science* vol. 8, no. 2, pp. 67–73. DOI: [10.1111/j.1467-9280.1997.tb00685.x](https://doi.org/10.1111/j.1467-9280.1997.tb00685.x).
- Ruby, MB and Heine, SJ (2011). "Meat, morals, and masculinity." *Appetite* vol. 56, no. 2, pp. 447–450. DOI: [10.1016/j.appet.2011.01.018](https://doi.org/10.1016/j.appet.2011.01.018).
- Rust, J and Schwitzgebel, E (2014). "The moral behavior of ethicists and the power of reason." In: *Advances in experimental moral psychology*. Ed. by H Sarkissian and JC Wright. London, Bloomsbury Academic, pp. 91–110.

- Ryan, L, Debenham, J, Brown, M, and Pascoe, W (2017). *Colonial frontier massacres in Eastern Australia 1788-1872*. URL: <http://hdl.handle.net/1959.13/1340762>.
- Schwitzgebel, E and Rust, J (2015). "The behavior of ethicists." In: *A companion to experimental philosophy*. Ed. by J Sytsma and W Buckwalter. Chichester, Blackwell Publishing, pp. 225–233.
- Scruton, R (1998). *Animal rights and wrongs*. 2nd. Metro Books, Glasgow.
- Singer, P (1975). *Animal liberation: the definitive classic of the animal movement*. 4th. Harper Collins, New York, 2009.
- Singer, P and Mason, J (2007). *The ethics of what we eat: why our food choices matter*. Rodale Press, United States of America.
- Smith, A (1774). *The theory of moral sentiments*. Cambridge University Press, Cambridge, England, 2002.
- Smith, A (1776). *Wealth of nations*. Generic NL Freebook Publisher, Hoboken, N.J. URL: [eBookCollection\(EBSCOhost\)](http://eBookCollection(EBSCOhost)).
- Smith, M (1994). *The moral problem*. Blackwell Publishers, Cambridge, Massachusetts.
- Stanovich, KE and West, RF (2000). "Individual differences in reasoning: implications for the rationality debate?" *The Behavioral and Brain Sciences* vol. 23, no. 5, pp. 645–665. DOI: [10.1017/S0140525X00003435](https://doi.org/10.1017/S0140525X00003435).
- Strawson, PF (1962). "Freedom and resentment." In: *Freedom and resentment and other essays*. 2nd. Milton Park, 2008, Routledge.
- Talbert, M (2006). "Contractualism and our duties to nonhuman animals." *Environmental Ethics* vol. 28, pp. 201–215.
- Thrasher, J and Mayne, L (2018). "Thinking about meat." *Unpublished manuscript*.
- Vanderschraaf, P (1995). "Convention as correlated equilibrium." *Erkenntnis* vol. 42, no. 1, pp. 65–87. URL: <http://www.jstor.org.ezproxy.lib.monash.edu.au/stable/20012605>.
- VanDeVeer, D (1979). "Of beasts, persons, and the original position." *The Monist* vol. 62, no. 3, pp. 368–377. URL: [url:http://www.jstor.org/stable/27902602](http://www.jstor.org/stable/27902602).
- Wheatley, T and Haidt, J (2005). "Hypnotic disgust makes moral judgments more severe." *Psychological Science* vol. 16, no. 10, pp. 780–784. DOI: [10.1111/j.1467-9280.2005.01614.x](https://doi.org/10.1111/j.1467-9280.2005.01614.x).
- Wicker, AW (1969). "Attitudes versus actions: the relationship of verbal and overt behavioral responses to attitude objects." *Journal of Social Issues* vol. 25, no. 4, pp. 41–78. DOI: [10.1111/j.1540-4560.1969.tb00619.x](https://doi.org/10.1111/j.1540-4560.1969.tb00619.x).
- Wiggins, D (1987). "A sensible subjectivism?" In: *Needs, values, and truth*. Oxford, Blackwell, pp. 184–214.
- Wilson, TD, Lindsey, S, and Schooler, TY (2000). "A model of dual attitudes." *Psychological Review* vol. 107, no. 1, pp. 101–126. DOI: [10.1037/0033-295X.107.1.101](https://doi.org/10.1037/0033-295X.107.1.101).

- Wolf, S (1982). "Moral saints." *The Journal of Philosophy* vol. 79, no. 8, pp. 419-439. URL: <http://www.jstor.org/stable/2026228>.
- Yadav, Y and Kumar, S (2006). "The food habits of a nation." *The Hindu*. URL: <http://www.thehindu.com/todays-paper/the-food-habits-of-a-nation/article3089973.ece>.
- Zaraska, M (2016). *Meathooked*. Basic Books, New York.

COLOPHON

This document was typeset using a modified version of the typographical look-and-feel classicthesis developed by André Miede and Ivo Pletikosić. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". classicthesis is available for both L^AT_EX and L^yX:

<https://bitbucket.org/amiede/classicthesis/>

Final Version as of April 3, 2018 (classicthesis).