

How We Value Animals

The Psychology of Speciesism

Lucius Caviola

Lincoln College

University of Oxford

July 2019

A Thesis

Submitted for the Degree of

Doctor of Philosophy

Word Count: 54,097

Abstract

In this thesis, consisting of three empirical projects (18 studies, total $N = 8,218$), I investigate how people morally value animals and why they value them less than humans. My central hypothesis is that people primarily value animals less than humans because of their mere species-membership, i.e. due to *speciesism*.

First, I develop and validate a Speciesism Scale, a psychological assessment instrument that measures the extent to which people value humans more than animals. I find that speciesism is a stable construct with large interpersonal differences. It shares psychological properties with other forms of prejudice and predicts behaviour directed towards animals.

Second, I experimentally investigate in more detail why people endorse speciesism. I develop a multi-factorial framework according to which people: a) believe that individuals should prioritise members of their own species over others (*Species-Relativism*); b) have a slight tendency to consider humans to be of superior value in an absolute sense (*Pro-Human Bias*); and c) have a specific bias against animals (*Anti-Animal Bias*). I also show that people value animals less, the lower their perceived mental capacity levels are (*Individual Mental Capacity View*)—a principle that people selectively apply to animals but not to humans.

Third, I experimentally investigate when people consider it permissible to harm animals or humans in cases where harming a few could save many. My findings support a model which I call *Multi-level Weighted Deontology*: deontological protections against harm are not absolute and get weaker the less people value the respective being, but they do not completely disappear. They are strongest for humans, followed by dogs, chimpanzees, pigs, and finally inanimate objects.

Overall, my research shows that speciesism is a pervasive psychological phenomenon that strongly impacts our social lives and moral thinking. This thesis is one of the first attempts to transition speciesism from moral philosophy into psychology.

Acknowledgments

I am grateful to my supervisors Nadira Faber and Julian Savulescu. They enabled me to pursue the research that I considered important even when the topics were unconventional. Nadira has taught me how to conduct rigorous psychological research, and how to transform my often abstract ideas into testable studies. Julian's deep understanding of the opportunities for cross-pollination between moral philosophy and experimental psychology has been an inspiration for me. Both have always been very supportive of me at a personal level.

I thank my close collaborators Guy Kahane and Stefan Schubert for their invaluable inputs. Guy's brilliant ideas have not just contributed very significantly to our joint projects but have also influenced my thinking about moral psychology. The countless hour-long discussions with Stefan have often been the place where insights arose and were further developed. For their constructive help with my research projects, I thank Fabienne Sandkühler, Jim Everett, and Elliot Teperman. For inspiring central ideas of this thesis and for providing useful feedback to the development of the Speciesism Scale, I thank Peter Singer, Oscar Horta, and Steven Wise.

For their comments, support, and discussions, I thank numerous friends, including David Althaus, Harri Besceci, Steven Bürgin, Natalie Cargill, Raffael Fasel, Ben Garfinkel, Jens Jäger, Simon John, Adriano Mannino, David Moss, Dominik Peters, Pablo Stafforini, Johannes Treutlein, Jonas Vollmer, and Matthieu Queloz. I am grateful to my parents, Carlotta and Hugo Caviola, and my sister, Silvia. Their encouragement to pursue whatever I was passionate about, even when it entailed quitting school to focus on computer programming, has led me via detours to write this thesis.

Finally, I thank those whose generous financial support enabled me to complete my studies and conduct my research: the Berrow Foundation at Lincoln College, the Swiss Study Foundation, the Janggen-Pöhn Foundation, the Oxford Martin School, and the Centre for Effective Altruism.

Contents

1. Introduction.....	5
The Speciesism Hypothesis	6
Is Speciesism a Form of Ingroup Favouritism?	9
Contribution to the Literature	12
Overview of Empirical Chapters	13
2. Speciesism as a Psychological Construct.....	17
Study 2.1. Scale Development	21
Study 2.2. Temporal Stability	30
Study 2.3. Convergent Validity	32
Study 2.4. Helping Via Donation Allocations	40
Study 2.5. Helping Via Time Investment and Food Choice	50
General Discussion	56
3. A Multi-Factorial Framework for Moral Anthropocentrism.....	61
Study 3.1. Individual Mental Capacity Level	64
Study 3.2. Species-Typical Mental Capacity Level	70
Study 3.3a. Species-Relativism and Pro-Human Bias	79
Study 3.3b. Species-Relativism and Pro-Human Bias	85
Study 3.4. Harming and Anti-Animal Bias	92
General Discussion	98
4. Deontological Constraints Against Harming Animals.....	108
Study 4.1. Humans vs. Animals	117
Study 4.2a. Different Types of Animals	122
Study 4.2b. Different Types of Animals	127

Study 4.3. Cross-Species Sacrifices	134
Study 4.4. Failed Attempt of Speciesism Reduction	140
Study 4.5. Between Objects and Humans	143
Study 4.6. Mental Capacities	149
Study 5.7. Time Investment and Donations	156
General Discussion	161
5. Concluding Discussion	171
Summary of Key Findings	171
The Relative Strength of Speciesism and Mental Capacities	173
Speciesism as a Form of Ingroup Favouritism	174
Differences in Moral Thinking About Humans and Animals	175
Limitations and Critical Reflection	177
Future Directions	179
Implications	180
Appendix.....	182
List of Figures	182
List of Tables	184
References	186

1. Introduction

Humans clearly prioritise themselves over animals¹. This shows in many different domains, across all cultures, and throughout history. Most conspicuously, humans kill and eat animals on a massive scale. Similarly, humans use some animals as experimental subjects to benefit humans. Yet other animals are kept as sources of entertainment or used as industrial equipment, or they are considered pests. Our legal system views even our closest non-human relatives, chimpanzees, as property and denies them basic rights such as the right to bodily integrity. I call this general phenomenon of morally valuing humans more than animals *moral anthropocentrism*. Moral anthropocentrism appears to be a strong and universal human tendency (Petrinovich, O’neill, & Jorgensen, 1993). In this thesis, I will investigate it psychologically. In particular, I will focus on the *speciesism hypothesis*: the hypothesis that people value humans more than animals because of their mere species-membership.

The fact that people value humans more than animals has also been demonstrated empirically using studies that rely on moral dilemmas in which the lives of humans are pitted against the lives of animals. These studies have found that people strongly prioritise humans over animals. In one study, for example, people generally prioritised the life of a human over the lives of animals, even when the human was a 75-year-old and the animals were the remaining members of an endangered species (Petrinovich et al., 1993). A recent large-scale survey assessing people’s judgments on how autonomous cars should act in moral dilemma situations showed that people exhibited a very strong preference to save humans over animals (Awad et al., 2018). Similarly, another study on dilemmas with autonomous cars found that pets were given fairly low moral weighting compared with most other categories such as man, woman, doctor, or child (Kim et al.,

¹ For simplicity purposes, I use the word “animals” instead of “non-human animals”. This includes all organisms that form the biological kingdom Animalia (i.e., multicellular eukaryotic organisms) but excludes humans.

2018). It has also been shown that people endorse actively harming animals to save humans (e.g. Topolski, Weaver, Martin, & McCoy, 2013).

While the existence of moral anthropocentrism is clear, its underlying psychological factors are poorly understood. Why do people value humans more than animals? Is it because animals are less intelligent than humans, that they cannot speak, are incapable of having complex thoughts, and cannot be held morally accountable for their actions? Or are there other reasons? Do we perhaps value animals less *just* because they are not humans?

The Speciesism Hypothesis

Psychologists have so far only paid relatively limited attention into understanding how people value animals. In contrast, philosophers have discussed human-animal relations for decades, if not centuries. While the primary focus of these philosophical discussions has been the normative issue of how we ought to value and treat animals, descriptive hypotheses about the psychological underpinnings of why people value animals less than humans often enter into them. My research is guided by these philosophical discussions, and in particular by the speciesism hypothesis.

In the philosophical literature, the term “speciesism” has been used to describe the view that individuals have different moral status merely because of their species-membership—i.e. independently of other differences, such as differing mental capacities (Horta, 2010; Ryder, 2006; Singer, 1975, 1979). The “speciesism hypothesis” is, in turn, the hypothesis that the root cause of moral anthropocentrism is speciesism (Singer, 1975). Even though popularised by philosophers, the speciesism hypothesis fundamentally involves a psychological hypothesis. In fact, the term speciesism was originally coined by the British psychologist Richard D. Ryder in Oxford in the early 1970s. (Ryder, 2010).

The main focus in the debate about speciesism has been on the view that considers humans more valuable than animals *just* because they are humans, which, if true, could partly explain moral anthropocentrism. However, speciesism could also manifest itself in the view that considers certain animals, such as dogs, more valuable than other animals with comparable mental and emotional capabilities, such as pigs. In both views, the moral status of certain individuals is determined based on their species-membership. While I will primarily focus on the former aspect of speciesism in this thesis, I will also, to some extent, investigate people's tendency to value certain animals more than others.

Philosophers often use the term speciesism in its normative sense of unjustified discrimination or prejudice, implying that people *should not* assign different moral values to individuals based solely on their species-membership, with analogies made with treating people different solely based upon their race (racism) or gender (sexism). However, this thesis is only concerned with speciesism in its psychological descriptive sense, namely that people *do*, as a matter of psychological fact, assign moral worth to individuals on the basis of species-membership.

One might ask whether it is really speciesism that best explains why we morally value humans more than animals (for philosophical criticism of speciesism as a concept see Diamond, 1978; Kagan, 2016; Williams, 2009). Could it not also be due to other properties that happen to correlate with species-membership? Perhaps the most common objections to the speciesism hypothesis, raised by philosophers and laypeople alike, are the following: rather than reflecting a speciesist bias, humans devalue animals because 1) animals are less sentient (i.e. able to experience feelings such as suffering) than humans; or 2) animals have much more basic cognitive capacities than humans (e.g. they cannot speak, do not have complex thoughts, cannot plan for the future, cannot make deliberate and autonomous judgments including moral evaluations, etc.). If any of these reasons could already fully explain why people value animals less than humans, the

speciesism hypothesis could be ruled out. However, when both of the above reasons are carefully analysed, they appear to be incomplete explanations of moral anthropocentrism (e.g. Horta, 2010).

First, the argument from sentience seems to fail because empirical data has shown that many animals—and in particular vertebrates, like pigs, dogs, and cows—are capable of suffering to an extent analogous to humans (Low et al., 2012). It is possible, however, that people systematically underappreciate the level of sentience in many animals. And as we will see in the next section, they do so to some extent (e.g. Bastian, Loughnan, Haslam, & Radke, 2012; Kozak, Marsh, & Wegner, 2006). However, it seems clear that people generally *do* perceive animals as sentient, which is made evident by their strong disapproval of unnecessary deliberate animal cruelty (Vaughn, 2009). People also do not typically doubt the sentience of human infants, despite the fact that the behaviour of infants offers weaker evidence for sentience than the behaviour of, for example, adult chimpanzees, which grieve over the death of relatives (Balter, 2010). Based on these preliminary considerations, it seems that moral anthropocentrism cannot convincingly be explained by the fact that people perceive animals to be less sentient.

Second, the argument from cognitive capacities seems to fail since it does not account for cases where people appear to value beings of one species less than beings of another species despite both having equal cognitive capacities or the less valued one even having more advanced cognitive capacities. Pigs, for example, have slightly more advanced cognitive capacities than dogs, and even pass a weak version of the mirror test, indicating some level of self-awareness (Broom, Sena, & Moynihan, 2009). It cannot, therefore, simply be cognitive capacities that determine treatment, because otherwise pigs would be treated as equal, and maybe even superior, to dogs. Of course, some people might morally value pigs less than dogs because they incorrectly assume that pigs are less intelligent than they actually are—a tendency which might be driven by motivated reasoning (Bastian et al., 2012). But such incorrect assumptions cannot explain clear cases of people valuing certain individuals less despite being fully aware that they are equally or more intelligent than

others. For example, it seems that most people would place less moral value on a chimpanzee than on a human with very severe cognitive disabilities even in cases where they know that the chimpanzee has more advanced cognitive capacities than said human (I tested this in Chapters 2 and 3). In fact, most people presumably grant equal moral status to severely cognitive disabled humans and healthy individuals despite their lack of cognitive and reciprocal abilities. Cognitive capacities, therefore, cannot fully account for the way we treat members of different species, because otherwise we would observe that people afford greater weight to the treatment of intelligent animals over severely cognitively disabled humans.

These considerations suggest that neither sentience nor cognitive capacities—or *mental capacities* to refer to both factors—can fully explain why people value humans more than animals. By this I do not mean that perceived mental capacities play no role at all in explaining moral anthropocentrism. In fact, in this thesis I will show that the perceived level of mental capacity of an animal does partially contribute to the moral status people attribute to that animal. Furthermore, in Chapter 3 I will explore more complex theories of how moral capacities could explain why people value animals less than humans. Overall, though, these considerations suggest that there must be a factor above and beyond mental capacities that explains why people value humans more than animals. I believe that this factor is species-membership itself. However, while the speciesism hypothesis may appear convincing at this point, it has not been empirically tested yet.

Is Speciesism a Form of Ingroup Favouritism?

While moral philosophers have debated possible reasons for prioritising humans over animals at great length, only limited psychological empirical research on the topic exists. Previous research has proposed that social psychological intergroup theories that describe human-human forms of outgroup mistreatment also apply to human-animal relations (e.g. Amiot & Bastian, 2015; Dhont, Hodson, & Leite, 2016; Plous, 1993). The underlying assumption is that animals are viewed

as members of outgroups similar to how people of other nations, cultures, or political affiliations are viewed as outgroup members (Amiot & Bastian, 2015). Relevant for the speciesism hypothesis is the question whether people classify animals as inferior outgroups simply because of their species-membership. While the previous research did not explicitly test this hypothesis, its findings are relevant for our understanding of how people value animals more generally. Note also that previous research sometimes used the term speciesism to describe the general phenomenon of valuing humans more than animals—which I would define as moral anthropocentrism—and not to describe the more specific speciesism hypothesis that suggests that this is *because* of species-membership². For simplicity reasons, however, I will use the term speciesism in the next sections when describing previous research. Throughout this thesis, I will provide evidence for the claim that not just moral anthropocentrism in general, but speciesism in particular, is a form of ingroup favouritism analogous to human-human forms of ingroup favouritisms.

Social Dominance Orientation

The first connection between speciesism and ingroup favouritism comes in the form of *social dominance orientation* (SDO)—an individual difference construct capturing support or opposition of group-based dominance and inequality amongst social groups (Pratto, Sidanius, Stallworth, & Malle, 1994). Differences in SDO predict prejudicial attitudes towards a variety of human social groups, including ethnic minorities, housewives, people with mental health difficulties, and people who are obese or perceived as unattractive (e.g. Duckitt & Sibley, 2007). Critically for a discussion of moral anthropocentrism, SDO also relates to how people feel about inequality between humans and animals. Recent studies suggest that people who believe in the superiority of humans over animals also believe in the superiority of some humans over others

² In Chapter 3, I will lay out a more systematic classification of the different concepts related to the moral value humans assign to animals.

(Costello & Hodson, 2010; Dhont & Hodson, 2014; Dhont, Hodson, Costello, & MacInnis, 2014). Accordingly, Dhont, Hodson and Leite (2016) have proposed the *Social Dominance Human-Animal Relations Model* (SD-HARM) whereby the same socio-ideological beliefs that legitimise hierarchies amongst human groups also legitimise hierarchies of humans over animals. In particular, SD-HARM assumes that human-human forms of prejudice like racism are psychologically connected to speciesism via an underlying tendency of SDO. In Chapter 2, I will provide empirical evidence supporting SD-HARM.

Dehumanization and Perceived Human-Animal Dissimilarity

The second connection between speciesism and ingroup favouritism comes from work on dehumanization of human outgroups, and specifically research suggesting that intergroup dehumanization is linked with beliefs in the superiority of humans over animals. Dehumanization is the psychological process by which certain outgroup members are seen as less “human” and therefore not worthy of full moral concern (Haslam, 2006). An example is the language of dehumanization, such as referring to black people as “apes”, Jews as “rats” and women as “bitches”, which is used to strip the victim of moral worth, as it is assumed that actual apes, rats, and dogs could not merit full moral consideration.

According to the *Interspecies Model of Prejudice* (ISMP; Costello & Hodson, 2014), the foundation of dehumanization is the belief that humans and animals are fundamentally different (i.e. human–animal divide). That is, only because people believe in a hierarchical divide between humans and animals considering certain human outgroup members as “inferior” dehumanization is even made possible (Costello & Hodson, 2010). Were people to view humans and animals as more similar, dehumanization would be ineffective. In support of this view, it has been demonstrated that when similarities between animals and humans are pointed out, not only negative attitudes towards animals is reduced but also moral concern for marginalised human outgroups is increased (Bastian et al., 2012; Costello & Hodson, 2010). One interpretation is that

when people perceive humans and animals as similar, they categorise them into a superordinate group that entails all humans as well as animals (Amiot & Bastian, 2015).

De-Mentalization

The dehumanization of human outgroup members often involves a process by which people reduce their attribution of mental states of others—a process referred to as *de-mentalization*. The same process has been observed with animals. It has been shown that a person's moral concern for animals is closely related to how much they believe animals can suffer (Waytz, Gray, Epley, & Wegner, 2010). Further, the denial of animals' capability to suffer can reduce moral concern for animals (Bastian et al., 2012; Kozak et al., 2006). De-mentalization has also been shown to drive the so-called *meat paradox*—the phenomenon that people do not want to hurt animals, yet continue to eat them (Bastian et al., 2012; Bratanova, Loughnan, & Bastian, 2011; Loughnan, Haslam, & Bastian, 2010). For example, people are more likely to de-mentalize animals if they had recently eaten meat (Loughnan et al., 2010) or if they anticipate eating meat soon (Bastian & Loughnan, 2017; Bastian et al., 2012).

Contribution to the Literature

In this thesis, I will provide an empirical investigation of the speciesism hypothesis. Moral anthropocentrism in general and speciesism in particular are neglected topics in psychology. For example, they are not mentioned in major social psychology text books (e.g. Hewstone, Stroebe, & Jonas, 2016) and in the past 75 years there were fewer than 50 papers published on these topics in academic psychological journals (based on a Web of Science search). This is surprising, considering that social and moral psychology investigate how people think about, interact with, and treat other living beings. Their focus, however, has been almost exclusively on relations with other humans. Only recently a new area of research has started to emerge that explicitly investigates

the psychology of how people view animals, which I presented further above³. My hope is to extend this previous research by giving both a more complete and systematic explanation of why people value humans more than animals. My research is guided by ideas developed in moral philosophy. I hope that this thesis will help to transition key ideas on the moral status of animals from philosophy into psychology.

Understanding how people value animals is important for a number of reasons. First, it contributes to a more complete psychological understanding of human social life and morality—one that also incorporates how we think, treat, and interact with animals. If people value animals differently than humans, does this mean that moral thinking about animals differs qualitatively from moral thinking about humans, or is the difference simply a matter in degree? And could insights into how people value animals also sharpen our understanding of how people value humans, in particular humans belonging to outgroups? Second, understanding how people value animals could have practical implications that may be of interest to philosophers and legal scholars. The way in which we value animals is a central and controversial aspect of human morality—one that shapes relations between humans and animals in numerous practical contexts.

Overview of Empirical Chapters

The thesis has three main aims. First, I will develop a Speciesism Scale—a theoretically driven and empirically validated measurement instrument to assess speciesism. Second, I will investigate the underlying factors of speciesism more deeply and explore what other factors, next to speciesism, contribute to moral anthropocentrism. Third, I will investigate a particularly relevant aspect of morality in the context of human-animal relations: deontological constraints against harming animals. I will now introduce the respective empirical chapters in more detail.

³ There also exists a field called anthrozoology—largely unknown to psychologists—which investigates human-animal relations more generally (cf. Herzog, 2010).

Speciesism as a Psychological Construct (Chapter 2)

In Chapter 2, I investigate speciesism as a descriptive psychological construct by developing a Speciesism Scale. This serves the purpose of demonstrating that speciesism is a real psychological construct suitable for empirical investigation, providing the basis for subsequent research. In five studies, using both general population samples online and student samples, I show that speciesism: a) is a measurable, stable construct with high interpersonal differences; that b) goes along with a cluster of other forms of prejudice; c) is able to predict real-world decision-making and behaviour. Study 2.1 presents the development and empirical validation of a theoretically driven Speciesism Scale, which captures individual differences in speciesist attitudes. Study 2.2 shows high test-retest reliability of the scale over a period of four weeks, suggesting that speciesism is stable over time. Study 2.3 found correlations between speciesism and racism, sexism, homophobia, and SDO, implying that similar mechanisms might underlie both speciesism and other well-researched forms of prejudice. The findings also support the hypothesis that speciesism is a manifestation of SDO. Further, Study 2.3 shows that speciesism positively correlates with political conservatism, system justification, right-wing authoritarianism, and negatively with actively open-minded thinking and empathic concern. Studies 2.4 and 2.5 demonstrate that speciesism is able to predict prosocial decision-making directed at animals (helping per charitable donations and per time investment) and behavioural food choices above and beyond existing related constructs. In sum, my studies show that participants morally valued individuals of certain species less than others even when beliefs about intelligence and sentience were accounted for.

This chapter is adapted from the following manuscript: Caviola, L., Everett, J. A. C., & Faber, N. S. (2019). The moral standing of animals: Towards a psychology of speciesism. *Journal of personality and social psychology*, 116(6), 1011.

A Multi-Factorial Framework for Moral Anthropocentrism (Chapter 3)

In Chapter 3, I investigate the underlying factors of moral anthropocentrism more deeply. Based on work in moral philosophy, I devised several hypotheses for moral anthropocentrism, which I tested in five experimental studies. I found that participants prioritised individual humans over individual animals even when the animals were described as having more advanced mental capacities than the humans (Study 3.1). Participants further prioritised humans over other species whose typical mental capacity level was described as on a par with that of humans (Study 3.2). However, while participants always prioritised humans, regardless of mental capacity, they generally did value individual non-human beings with more advanced mental capacities more (Individual Mental Capacity View). To large extent, however, participants prioritised humans over animals purely because of their species-membership, i.e. speciesism. Most participants believed that all beings should prioritise members of their own species over members of other species (Species-Relativism; Studies 3.3a-b). In addition, I found a slight tendency to consider humans to be of superior value in an absolute sense (Pro-Human Bias). Finally and unexpectedly, I found a specific bias against animals that did not manifest towards hypothetical non-humans with equivalent mental capacities (Anti-Animal Bias; Study 3.4). Overall, this chapter demonstrates that there are multiple factors underlying moral anthropocentrism as well as speciesism, and offers a framework to study those in-depth.

This chapter is adapted from the following manuscript: Caviola, L., Schubert, S., Kahane, G., & Faber, N.S. (2019). Why People Prioritize Humans Over Animals: A Framework for Moral Anthropocentrism. *Manuscript in preparation.*

Deontological Constraints Against Harming Animals (Chapter 4)

In Chapter 4, I investigate whether the moral rules that people use to decide whether harm to humans is permissible are also applied in the context of harm to animals. In eight experimental studies I show that participants considered harming a few animals to save a greater number of

animals to be more morally permissible than harming a few humans to save a greater number of humans (e.g. Study 4.1). The deontological rules that constrain such trade-offs in the human case are thus not applied in the same way in the animal case. However, participants also did not decide whether harm to animals is morally permissible by engaging in strict cost-benefit analysis (Study 4.3). Instead, the model that best captures participants' moral judgements is what I call *Multi-level Weighted Deontology*: deontological constraints are not absolute and get weaker the lower the individual's level of moral status, but they do not completely disappear. They were strongest for humans, followed by dogs, chimpanzees, pigs (Study 4.2a-b) and finally inanimate objects (Study 4.5). Species-membership itself strongly influenced moral judgments, while the perceived level of an individual's cognitive capacity only made a small difference and sentience made none (Study 4.6). Finally, participants were more willing to invest time and money in campaigns that opposed medical research conducted on humans to benefit humans than research on pigs to benefit pigs (Study 4.7). My findings suggest that there is only a difference in degree between the way people think about harm to humans and harm to animals, and that people do not have an all-or-nothing understanding of deontological rules.

This chapter is adapted from the following manuscript: Caviola, L. Kahane, G., Everett, J.A.C., Teperman, E., Savulescu, J., & Faber, N.S. (2019). Utilitarianism for Animals, Kantianism for People? Speciesism in Sacrificial Moral Dilemmas. *Manuscript under review*.

2. Speciesism as a Psychological Construct

In this chapter, I investigate speciesism as a psychological construct. Specifically, I am interested in the empirical truth of the psychological claims implicit in the philosophical discussion of speciesism: first, the primary claim that people assign moral worth to individuals on the basis of species-membership alone; and second, the claim that speciesism is a form of prejudice analogous to other prejudicial attitudes. Philosophers have debated these claims, but relatively little empirical work has been conducted to test whether these claims are, as a matter of fact, true. Do people, in actuality, assign moral worth to individuals on the basis of species-membership; and are these speciesist attitudes connected to other prejudicial attitudes? By rigorously examining this, it becomes possible to bring the concept of speciesism into the study of intergroup relations and prejudice more generally, thus providing new insights and directions for research for both topics.

My preliminary aim, upon which all later aims of this chapter depend, was (Aim I) to develop a reliable and valid scale to measure speciesism as a psychological construct (Studies 2.1 and 2.2). Once this was achieved, I aimed to look at (Aim II) the extent to which speciesism can psychologically be considered a form of a prejudice (Study 2.3), by testing whether (Aim IIa) speciesism correlates with other forms of prejudice, and whether (Aim IIb) speciesism, like other forms of prejudice, is driven by socio-ideological factors such as social dominance orientation that maintain other forms of intergroup conflict. To the extent that speciesism is a form of prejudice it should be correlated with prototypical examples of prejudice and be driven by the same kind of processes that drive other prejudices. Moreover, I aimed to explore the relationship between speciesism with empathic concern and actively open-minded thinking (Aim IIc). Finally, I aimed to look at (Aim III) whether speciesism predicts behaviour, such as the degree of help

directed towards individuals of different species or the likelihood of choosing certain food products over others (Studies 2.4 and 2.5).

To what extent is it reasonable, on empirical grounds, to view speciesism as a form of prejudice? It is here that psychology has a real and important role to play. If it can be shown that speciesism is psychologically related to other forms of prejudice, the philosophical case for it can be strengthened. One typical definition of prejudice is that it refers to “any attitude, emotion, or behaviour toward members of a group, which directly or indirectly implies some negativity or antipathy toward that group” (Brown et al., 2010). Speciesism seems to fit that definition as it involves negative beliefs, emotions, and behaviour directed towards others based on species-membership. Most people believe that pigs matter less than dogs (attitude), feel disgust towards rats but love towards cats (emotion), and accordingly treat pigs and rats much worse than they treat dogs and cats (behaviour). Moreover, in addition to fitting the definition of prejudice, I aim to show that speciesism shares properties and underlying psychological mechanisms with other phenomena referred to as prejudice.

Measuring Speciesism

To study speciesism, a convenient way of measuring speciesist attitudes is needed. There are currently two established scales that attempt to capture general attitudes towards animals: Herzog, Betchart, and Pittman’s (1991) *Animal Attitude Scale* and Wuensch, Jenkins, and Poteat’s (2002) *Animal Rights Scale*. Both scales capture important aspects of people’s attitudes towards animals. Furthermore, Dhont et al. (2014, 2016) and Piazza et al. (2015) have developed ad hoc measurement instruments for speciesism, which are partly based on items of Herzog et al. (1991).

For several reasons, however, I argue that a new scale is needed in order to examine the psychology of speciesism specifically with sufficient precision. While these existing measures have been used to good effect to advance our understanding of human-animal relationships, they suffer from theoretical limitations and were produced without statistical validation and established scale

development procedures. A comprehensive validated measuring tool for speciesism, therefore, remains unavailable in the existing literature.

First, existing scales do not explicitly capture speciesism, according to its precise definition: attributing moral status to an individual solely on the basis of their species-membership. Instead, existing scales capture related, yet distinct, concepts such as general attitudes towards animals or views on animal rights.

Second, some items in existing scales confound empirical and normative issues. Consider the following item by Wuensch et al. (2002): “Most cosmetics research done on animals is unnecessary and invalid”. This item is limited as it confounds the normative belief that animals should not be subjected to suffering with the empirical belief in the efficacy of scientific testing. There is no way to accurately respond to the item if, for example, you believe the research is scientifically valid but morally abhorrent, or if you believe the research is morally acceptable but scientifically useless. Or consider another item on the scale, which reads: “There are plenty of viable alternatives to the use of animals in biomedical and behavioural research” (Wuensch et al., 2002). Again, this item is flawed as it relates to an empirical fact, not a belief about the moral status of animals. A coherent anti-speciesist could believe that there are few viable alternatives to the use of animals in research while maintaining testing on animals is morally wrong.

Third, a common difficulty in capturing speciesism is that people who endorse anti-speciesism can come to different conclusions about certain practices depending on the philosophical position they hold, such as consequentialism and deontology. Consequentialism, including utilitarianism, is the moral view that the rightness of an action depends only on its consequences (Bentham, 1780; Mill, 1861). Deontology is the moral view that certain actions are forbidden irrespective of the consequences (Kant, 1785). A consequentialist anti-speciesist might, for example, agree to harm animals (or humans) if this would result in better overall outcome for all sentient beings (Kahane et al., 2018; Kahane, Everett, Earp, Farias, & Savulescu,

2015). A deontological anti-speciesist, however, would refuse to harm animals (or humans) even if the outcome would be better for everyone. Given this divergence, unless carefully devised, items that, for example, assess people's attitudes on using animals for medical testing can be unhelpful (e.g. some items by Herzog et al., 1991). This is because they might distinguish between a consequentialist anti-speciesist (who may endorse research using animals if it results in the development of a drug that saves many lives in the long run), and a deontological anti-speciesist (who may reject research using animals because it violates a general rule of not harming others; cf. *Cross-Species Deontology* in Chapter 4). Such a distinction, however, is a factor that should not be captured by a speciesism scale as, ultimately, the scale must be able to identify speciesism without being confounded by unrelated specifics of the underlying moral position a person may hold.

Requirements for a Speciesism Scale

In order to ensure that my scale captures the precise philosophical meaning of speciesism, I outlined the following requirements to be met:

First, the initial set of items (subjected to exploratory factor analysis) should capture speciesism both exclusively and exhaustively. I assume that speciesism will manifest itself both in general beliefs about the moral inferiority of certain species and in the endorsement of concrete practices involving the use of animals. All major manifestations of speciesism must be covered, and core manifestations of speciesism must receive appropriate representation, such as the use of animals for entertainment, food, and medical experiments.

In order to capture all relevant manifestations of speciesism the item pool should consist of both abstract-general and concrete-empirical items. There are benefits and drawbacks to including concrete items. On the one hand, benefits of concrete items include the fact that people's attitudes are likely to be inconsistent and limited by self-serving biases, and concrete items can expose this inconsistency. For example, people might agree with an abstract item such

as “species-membership is not a morally relevant criterion” but disagree with a concrete item such as “animals should not be hunted for sport”. Another advantage of concrete items is that they may be easier for laypeople to understand and respond to. As this scale is primarily for use by laypeople and not philosophers, ease of comprehension is an important factor. On the other hand, disadvantages of concrete items include the fact that they inevitably create empirical confounders. Take a person’s views on animals performing in the circus, for example. To some extent, the issue poses the empirical question of how far circus animals suffer. People might have different experiences of circuses, which might cause them to have different views on the level of suffering circuses cause for animals. These beliefs about circus conditions might be entirely independent of their moral values. Given these arguments about concrete items, I aimed to strike an appropriate balance between philosophically rigorous items and concrete items in the final scale.

Second, items should avoid normative confounders. In particular, it is important that items do not prompt different responses from deontological and consequentialist anti-speciesists. As I have discussed above, there is the potential for consequentialist and deontological anti-speciesists to be divided on a number of issues related to the treatment or rights of animals. Therefore, when I referred to an empirical situation in my items, I ensured the situation was such that most consequentialist and deontological (anti-)speciesists would reach the same conclusion.

Study 2.1. Scale Development

In Study 2.1, I began my process of developing the Speciesism Scale (Aim I) by formulating and testing a pool of theoretically-validated items to measure speciesism.

Method

Item Selection. To formulate a list of items, I used a combination of partially modified versions of items from existing scales (Herzog et al., 1991; Wuensch et al., 2002) and entirely new items. I then sent these proposed items to renowned experts from a number of disciplines, including the philosophers Peter Singer and Oscar Horta, and the legal scholar Steven Wise¹. After incorporating feedback from these experts and excluding items that did not meet my criteria, I was left with a 27-item pool (*Table 2.1*). In line with my prerequisites for the scale, these items included both concrete items (e.g. “It is morally acceptable for cattle and pigs to be raised for human consumption”) and abstract items (e.g. “Some beings are morally more important than others just because they belong to a certain species.”); and items tapping both the belief in the superiority of humans over animals (“Morally, animals always count for less than humans”) and the superiority of certain animals over other animals (“Pigs should be taken care of by humans just like dogs are”: reverse-scored).

Ethics Statement. For all studies in this chapter, my institution’s ethical guidelines were followed, and the research was approved through University of Oxford’s Central University Research Ethics Committee, with the reference number MSD-IDREC-C1-2014-133.

Participants and Procedure. 1,122 US American participants took part in the study online via Amazon Mechanical Turk (MTurk) and received a payment of \$0.50 for their participation. Five participants were excluded for not completing the study, and nine because they failed a simple attention check embedded in the survey where people were asked to select a certain scale point to confirm they were paying attention. This left a final sample of 1,108 participants (457 female; $M_{age} = 33$, $SD = 11.56$). In contrast to experimental studies there are no straightforward and commonly accepted techniques to determine sample size for factor analyses (Mundfrom, Shaw, & Ke, 2005). One approach is to rely on the absolute sample size. Estimations by Comrey and Lee (2013) suggest that a sample size of 300 is adequate, and that 1,000 or more is

excellent. In general, the literature agrees that in exploratory factor analyses the higher the sample size the better (Osborne, Costello, & Kellow, 2008). A different approach is to aim for a subject-to-item ratio of at least 5:1, better 10:1 or higher (Everitt, 1975; Nunnally, Bernstein, & Berge, 1967). My final sample of 1,108 therefore represents a size that is more than adequate, and with 27 items, gave us an excellent final subject-to-item ratio of 41:1.

Items were presented in randomised order and participants were asked to indicate to which extent they agreed with the statements on a Likert scale from 1 (*strongly disagree*), over 4 (*neither agree nor disagree*) to 7 (*strongly agree*). If not further specified this response scale was used for all established measures throughout my five studies.

Results

Exploratory Factor Analysis. I first conducted Exploratory Factor Analysis (EFA) using the principal axis factors extraction method to determine the factor structure of the 27 items. Oblique direct oblimin rotation was chosen as I expected that underlying factors would correlate. The Kaiser-Meyer-Olkin (KMO) measure of sampling adequacy was .94, which is above the recommended value of .6, indicating that individual items shared enough common variance for such an analysis. Bartlett's test of sphericity was significant: $\chi^2(351) = 12,748.82, p < .001$. The first factor explained 34.04% of the variance (eigenvalue = 9.20), the second factor 9.39% of the variance (eigenvalue = 2.54), the third factor 5.30% of the variance (eigenvalue = 1.43), the fourth factor 5.08% of the variance (eigenvalue = 1.37), and the fifth factor 3.83% of the variance (eigenvalue = 1.04).

The scree plot displayed a clear inflexion point, which justified retaining two factors. As such, I ran a further analysis, specified to extract two factors. The first extracted factor seemed to capture my intended construct of speciesism including both abstract items and more concrete items capturing attitudes towards animal exploitation and animal rights. The second factor captured ethical vegetarianism (e.g. "It is morally wrong to eat fish").

Based on theoretical and statistical (see CFA below) considerations, I decided to retain the first factor only. It best captured the “pure” theoretical construct of speciesism, in isolation from other constructs. In this context, it is notable that factor analysis extracted speciesism and ethical vegetarianism as separate constructs. This is evidence that the two constructs are psychologically distinct. The speciesism factor and the ethical vegetarianism factor correlated negatively with each other, $r = -.35, p < .001$.

Items for the further development of the scale were selected or excluded on the basis of theoretical and statistical considerations. In particular, I focused on including items with high factor loadings while ensuring that theoretically relevant aspects of speciesism were captured by the set of items. I prioritised items that did not include empirical assumptions about the intelligence or suffering of animals to avoid confounding assumptions (e.g. items three or eight). My set of selected items contained ten items (Table 2.1).

Table 2.1. Factor loadings from the second EFA (Study 2.1).

Item	Item Label	Factor 1	Factor 2
1*	Humans have the right to use animals however they want to.	.742	
2*	It is morally acceptable to keep animals in circuses for human entertainment.	.701	
3	It is morally acceptable to kill animals, because they are less intelligent than humans.	.682	
4*	It is morally acceptable to hunt animals for sport.	.665	
5*	Morally, animals always count for less than humans.	.638	
6*	It is morally acceptable to do cosmetic research on animals.	.631	
7*	Some beings are morally more important than others just because they belong to a certain species.	.602	
8	A pig has a lower capacity to suffer than a human baby.	.598	

9*	It is morally acceptable to perform medical experiments on animals that we would not perform on any human.	.597
10	A pig has a lower capacity to suffer than a dog.	.585
11	It is morally acceptable to kill animals even it is not necessary for our own survival.	.582
12*	Basic rights that are enjoyed by humans – like protection from harm – should also be granted to animals.	-.582
13	Reducing pain and suffering in animals is morally equally important as reducing pain and suffering in humans.	-.577
14	It is morally acceptable to breed animals just to produce leather out of their skin.	.560
15*	It is morally acceptable to trade animals like possessions.	.545
16	If an elephant hurts itself in nature, it is morally acceptable not to help even if it were morally required to help a human in the same situation.	.532
17*	Chimpanzees should have basic legal rights such as a right to life or a prohibition of torture.	-.527
18	It is morally acceptable to kill animals that destroy human property, for example, rats, mice, or pigeons.	.502
19	It is morally worse to kill any human than to kill a chimpanzee.	.485
20	Factory farming of animals is morally wrong.	-.397
21	Pigs should be taken care of by humans just like dogs are.	-.373
22	Faced with a decision of killing one human embryo or one pig, we should always kill the pig.	.365
23†	It is morally wrong to eat fish.	.820
24†	It is morally wrong to consume milk and eggs.	.743

25†	It is morally required to become vegetarian in an effort to save animals.	.648
26†	It is morally acceptable for cattle and pigs to be raised for human consumption.	-.529
27†	It is morally acceptable to hunt wild animals for food.	-.528

Notes. * included in CFA's first factor, † included in CFA's second factor

Confirmatory Factor Analysis. In order to examine the model fit of the ten extracted items I performed a Confirmatory Factor Analysis (CFA) relying on maximum likelihood estimation in R using the *lavaan* package (version 0.5).

I used a combination of fit indices to judge the model fit. My primary model fit index was the Comparative Fit Index (CFI) score, which indicates how well the data fits the target model compared to an independent model that assumes uncorrelated variables. CFI is frequently reported and, in comparison to other fit indices, not unduly influenced by sample size (Fan, Thompson, & Wang, 1999). Based on recommendations in the literature I considered a model to have an acceptable fit if its CFI score was .93 or higher (Mendl, Held, & Byrne, 2010). Due to the robustness of the CFI I decided to reject any model that would not meet the required CFI standard.

For models with an acceptable fit according to CFI I furthermore considered the Tucker-Lewis Index (TLI), which is an incremental fit index and must lie above .90, and ideally above .95 (Hu & Bentler, 1999; Marsh, Balla, & McDonald, 1988). I also applied the more traditional Chi-square test. However, it must be noted that the Chi-square test for CFA is considered to be an unreliable fit index since it is very sensitive to sample size and will usually result in significant outcomes for sample sizes larger than 200 (Steiger, 2007). Finally, I considered the Standardised Root Mean Square Residual (SRMR), which represents the standardised discrepancy between the

predicted and observed correlation and must lie below .08 to justify adequate model fit (Hu & Bentler, 1999).

A first CFA revealed an unsatisfactory model fit with a CFI of 0.88, which lies below the required standard (TLI = .85; SRMR = .06). Results showed that unexplained variances correlated highly among each other. This implied that certain items shared variance among each other that was not captured by general speciesism. For example, items two and four both captured the use of animals for entertainment and as such tapped into a distinct sub-form of speciesism. Similar redundancies were found between items five and seven that were both highly abstract without tapping into real-world examples, items six and nine both tapped into the domain of research use with animals, and items twelve and seventeen both captured attitudes towards animal rights. Given these redundancies, I chose to exclude items four, six, seven, and twelve from the final scale, as they did not serve any additional explanatory purpose.

A second CFA with the remaining six items was then conducted. The CFI was .98, which suggests adequate model fit. TLI was .96, which indicates excellent fit. SRMR was .07, which is in the range of acceptable values. In contrast, to the already mentioned fit indices, the Chi-square test yielded a poor fit, $\chi^2(9) = 52.87, p < .001$. However, as mentioned the Chi-square test is considered to be an unreliable fit index for which reason I did not strongly weight its result but still report it for completeness sake. As such, CFA indicated that the model of the six-item scale was an appropriate fit. Furthermore, a Kolmogorov-Smirnov test for normality suggested that the scores were normally distributed, $D(1108) = .02, p = .20$.

Finally, I conducted a third CFA in which I tested the model fit of a two-factor model, which includes both the speciesism and ethical vegetarianism factor. The model included the six speciesism items and five ethical vegetarianism items that loaded onto the second extracted factor (*Table 2.1*). The model fit, however, was unsatisfactory (CFI = .88; TLI = .85; SRMR = .08). Only after removing items 26 and 27 the model fit became satisfactory (CFI = .96; TLI = .94; SRMR =

.05). However, due to its low item count and its non-normally distributed scores (Kolmogorov-Smirnov test for normality: $D(1108) = .11, p < .001$) I do not recommend using the ethical vegetarianism factor as a measurement instrument.

Confirmatory Factor Analysis on a Separate Dataset. In order to confirm my six-item one-factor solution I collected a new set of data with which to test my model. 200 US American participants took part in the study online via MTurk and received \$0.70 for taking part. Four participants did not complete and nine were excluded for failing the attention check question. This left a final sample of 187 US American participants (89 female; $M_{age} = 35.73, SD = 10.36$). This sample size is again adequate according to the rule of thumb of aiming for a subject-to-item ratio of at least 1:10 (Everitt, 1975; Nunnally et al., 1967) as in my case the ratio was 31:1. Furthermore, my sample size nearly equals Jackson’s (2001) recommendation to recruit a minimum of 200 participants for CFA relying on maximum likelihood estimation. Participants responded to the speciesism items and demographic questions. This new data set provided very strong evidence for the six-item model, where the CFI was .99, TLI was .99, the SRMR was .02, and even the Chi-square test yielded an acceptable fit, $\chi^2(9) = 11.87, p = .22$.

The Speciesism Scale. Based on a series of EFA and CFAs, I arrived at the final items to form a Speciesism Scale (

Table 2.2). The Speciesism Scale consists of six items all loading onto a single factor. Speciesism scores were normally distributed across the sample with a mean of 3.64 ($SD = 1.25$) where the minimum was 1 and the maximum 7. Reliability analyses yielded that the six-items scale had high internal consistency, with a Cronbach’s alpha of .81.

Table 2.2. Speciesism Scale

Item	Item Label
------	------------

- 1 Morally, animals always count for less than humans.
- 2 Humans have the right to use animals however they want to.
- 3 It is morally acceptable to keep animals in circuses for human entertainment.
- 4 It is morally acceptable to trade animals like possessions.
- 5 Chimpanzees should have basic legal rights such as a right to life or a prohibition of torture. (r)
- 6 It is morally acceptable to perform medical experiments on animals that we would not perform on any human.

Notes. Answers on 7-point scale from "Strongly disagree" to "Strongly agree". (r) indicates reversed scoring.

Discussion

The goal of Study 2.1 was to develop a theoretically informed and empirically sound *Speciesism Scale*. A combination of EFA, CFA, and reliability analysis supported the development of a one-dimensional scale. All six items strongly loaded onto a single speciesism factor, which had good model fit and high internal consistency, and this model was confirmed by an additional CFA conducted on a separate dataset. Most importantly, the scale fulfils all requirements that were stipulated at the beginning of the study: all items explicitly capture speciesism and the scale encompasses crucial aspects of the theoretical concept; experts in relevant fields validated all items; the scale consists of both abstract and empirical items; and does not contain items eliciting empirical and normative confounding factors.

It is interesting that speciesism and ethical vegetarianism were—despite being strongly correlated—psychologically distinct factors. At first blush, this result might be surprising, as one might expect that endorsement of anti-speciesism would consistently result in endorsement of ethical vegetarianism. However, this finding is consistent with previous research on the meat paradox (Bastian & Loughnan, 2017; Bratanova et al., 2011). People might endorse anti-speciesism

in the abstract or in domains where it does not conflict with their personal preferences, but they employ specific beliefs and practices in the context of food (i.e. carnism; Monteiro, Pfeiler, Patterson, & Milburn, 2017; Piazza et al., 2015).

Study 2.2. Temporal Stability

I have hypothesised that speciesism is a psychological form of prejudice analogous to other psychological forms of prejudice such as racism or sexism. If this is the case speciesism should—like other forms of prejudice—be a relatively stable construct that persists over time. Of course, like any other form of prejudice, the extent to which a person holds speciesist attitudes can fluctuate depending on the situation and other conditional factors. Ultimately, however, the extent to which a person holds speciesist attitudes should stay relatively stable over time because otherwise speciesism could not be considered a psychological prejudice analogous to racism or sexism, but a short-term belief connected to spontaneous emotional reactions. In order to investigate the temporal stability of speciesist attitudes, I examined the test-retest reliability of the Speciesism Scale over a period of four weeks. To the extent that speciesism does represent a stable prejudice, scores at the two times should be highly correlated.

Methods

The study consisted of two stages: a first stage in which participants completed the speciesism scale, and then a second stage four weeks later in which these same participants were invited to again complete the scale. In the first stage, 685 participants took part via MTurk and received a payment of \$0.50 for their participation. Eight participants were excluded for failing the attention check question, leaving a final sample of 677 US American participants (305 female; $M_{age} = 34.37$, $SD = 10.94$). I conservatively assumed a response rate of around 30% for the second stage of the study and an expected effect size (Pearson correlation coefficient) between 0.6 and 0.8. Based on sample size calculations for test-retest analyses in the literature (Shoukri et al., 2004)

I decided to aim for at least 200 responding participants in the second stage and as such aimed to recruit 680 in the first stage (given the expected 30% response rate). Participants completed the Speciesism Scale but were not informed about the follow-up study. Other filler measures were included to distract from the speciesism items.

In the second stage four weeks later, all 677 participants were contacted again with information about the retest study. 333 US American participants (164 female; $M_{age} = 36.82$, $SD = 11.51$) completed the retest study (giving a higher-than-expected response rate of 49%), and all participants again received \$0.50 payment for taking part. No participants were excluded at this stage, and this final sample of 333 is more than adequate, exceeding minimum recommendations for test-retest analyses (Shoukri et al., 2004).

Results and Discussion

Internal consistency was high for both stages of the study with a Cronbach alpha of .89 for the first stage and .90 for the second stage of the study, and—critically—a retest analysis with the 333 cases revealed a test-retest correlation coefficient of $r = .88$, $p < .001$. That is, participants' scores on the Speciesism Scale were very highly correlated with their scores on the same scale four weeks later. The high test-retest reliability is an important part of establishing speciesism as a temporally stable psychological prejudice and demonstrates that speciesism—similar to racism or sexism—is not just a short-term belief or emotional reaction, but a stable view that persists over time.

Although my primary purpose for this study was to consider test-retest reliability, given that for the first sample I had a sample size of 677, I decided to subject the resulting speciesism ratings to an additional CFA to investigate if the six-item one-factor structure revealed in the two independent samples of Study 2.1 holds a third confirmatory test. Results confirmed once again that the six-item model had excellent fit (CFI = 0.98, TLI = 0.97, SRMR = 0.02).

Overall, in Study 2.2 I provided further psychometric support for my Speciesism Scale, finding excellent fit for my model in an independent data sample and finding good test-retest reliability over time. The results from Study 2.2 cannot show that speciesism is a result of similar psychological processes to other forms of prejudice, only that speciesism can in fact be considered a stable and measurable attitude like racism and sexism. In Study 2.3 I turned to look at how similar speciesism is to other forms of prejudice in terms of content.

Study 2.3. Convergent Validity

In Study 2.1 I developed a new and reliable scale to measure speciesism, and in Study 2.2, I confirmed that speciesism scores were persistent over time. Combined, the results of Studies 2.1 and 2.2 highlight that speciesism is a stable psychological construct. Having met my first aim—to develop a reliable and valid scale to measure speciesism—I next turned to my second aim: to better understand the nature of speciesism by looking at what other attitudes speciesism is associated with, and which more general psychological orientations drive such speciesist attitudes. By doing so I sought to establish convergent validity of the Speciesism Scale.

In Study 2.3, equipped with my new scale, I explored whether speciesism might be considered a form of prejudice by shedding light on the relationship of speciesism with other psychological constructs. I took a two-pronged approach: first, I tested whether speciesism correlates with other forms of prejudice (Aim IIa); and second, I looked at whether speciesism, like other forms of prejudice, is driven by socio-ideological factors such as social dominance orientation or right-wing authoritarianism that maintain other forms of intergroup conflict (Aim IIb).

My first aim—Aim IIa—was drawn from the personality approach to prejudice, which suggests that prejudice is typically a generalised phenomenon: a person who is high on ethnic prejudice will also be high on gender-based prejudice, and so on (Allport, Clark, & Pettigrew,

1954). For example, prejudice toward various targets tend to be significantly correlated (Akrami, Ekehammar, & Bergh, 2011), and factor analyses yield a generalised prejudice factor explaining 50% to 60% of the variance in different forms of prejudice (Ekehammar & Akrami, 2003). If speciesism is indeed to be usefully considered a form of prejudice, people who hold stronger speciesist attitudes should also exhibit other prejudicial attitudes.

My second aim—Aim IIb—is a result of the reasoning, as suggested by the SD-HARM model (Dhont et al., 2016), that if speciesism is a form of prejudice it should share psychological roots with other forms of prejudice. In this way, I hoped both to further validate the Speciesism Scale by showing it to be correlated with social dominance orientation and related socio-ideological constructs, as well as providing independent support for the claims of the SD-HARM model.

As a third and more exploratory aim—Aim IIc—, I aimed to investigate the relationship between speciesism with empathic concern and actively open-minded thinking. A previous study has already identified a relation between empathy levels and attitudes to animals (Taylor & Signal, 2005) and as such I hypothesised that people higher in empathic concern would care more about the suffering of animals and subsequently endorse less speciesist attitudes. Similarly, because speciesist attitudes predominate in society, I predicted that actively open-minded people, people who are more willing to change their beliefs (Baron, 2000) and think beyond the currently accepted norms, are more likely to endorse anti-speciesism.

Method and Results

257 US American participants took part in the study online via MTurk and received \$1 payment for their participation. Fifteen participants were excluded for failing an attention check, leaving a final sample of 242 people (110 female; $M_{age} = 36.33$, $SD = 11.88$). With that, my sample size met the recommendations of Comrey and Lee (2013) to recruit at least 200-300 participants for scale validation. Furthermore, a review of over a hundred scale validation studies revealed that the median sample size of such studies was 121, and so with a final sample of 242 I was confident

that my sample size was more than adequate (Anthoine, Moret, Regnault, Sébille, & Hardouin, 2014).

In the main part of the study, participants were asked to complete the Speciesism Scale and a number of separate scales presented in a random order, which are discussed in turn along with their results below. At the end of the study, participants responded to demographic questions including age, gender, education (six-step continuous scale from “less than high school degree” to “graduate degree”), income (10-step continuous scale from “under \$5,000 per year” to “over \$100,000 per year”), and whether they are vegetarian (yes or no). Median education level was “attended college” ($M_{education} = 4.22$, $SD = 1.29$) and median annual income level was “\$25,001-\$35,000” ($M_{income} = 5.22$, $SD = 2.34$).

I note that two items from the original item pool that in Study 2.1 were shown not to load on the main speciesism factor were included in this (and all following) studies. To ensure that my results for the main Speciesism Scale hold over and beyond the inclusion of these two items, I explored whether results changed when these two items were added to the scale. As it did not, this issue is further ignored. Due to the multiple correlation analyses I conducted with speciesism I relied on a Bonferroni adjusted alpha level of .006 per test (.05/9) as an indicator for statistical significance (see *Table 2.3* for all correlation coefficients).

Prejudicial Attitudes. My first set of measures was intended to address Aim 2a: to what extent does speciesism correlate with other, prototypical forms of prejudice? To the extent that speciesism is—as philosophers have argued it to be—a form of prejudice, it should be correlated with other forms of prejudice. Specifically, I looked at three prototypical forms of prejudice: prejudice based on ethnic background or race (racism), prejudice based on gender (sexism), and prejudice based on sexual orientation (homophobia). A growing body of research has suggested that speciesism is associated with prejudice, but most of this has used simple feeling thermometer

type questions (Dhont et al., 2014, 2016). In my study, therefore I drew upon this work but relied on widely used and empirically validated scales.

First, to tap racism, I used the Modern Racism Scale (McConahay, 1986). This consists of 7 items and includes items such as “Blacks are getting too demanding in their push for equal rights” (internal consistency: $a = .75$). Second, to look at sexism I used the Modern Sexism Scale (Swim, Aikin, Hall, & Hunter, 1995). This consists of 8 items and included items such as “Women often miss out on good jobs due to sexual discrimination” ($a = .92$). Third, to look at homophobia, I used the revised short version of the Attitudes Toward Lesbians And Gay Men Scale (Herek, 1998). This consists of 10 items including “The idea of homosexual marriages seems ridiculous to me” ($a = .97$). To the extent that speciesism is a form of prejudice analogous to other kinds, speciesist attitudes should be associated with increased ethnic, gender, and sexuality-based prejudice, just like these types of prejudice are typically associated with one another (Pratto et al., 1994). Indeed, confirming the contention that speciesism is correlated with other forms of prejudice, I found significant positive correlations of speciesism with racism ($r = .32, p < .001$), sexism ($r = .41, p < .001$), and homophobia ($r = .17, p < .001$).

Socio-Ideological Beliefs. My second set of measures were intended to address Aim 2b: to explore whether speciesism, like other forms of prejudice, is driven by socio-ideological factors. Four such socio-ideological factors stand out in previous research on prejudice and intergroup conflict: Social Dominance Orientation (SDO; Pratto et al., 1994); political conservatism; system-justification (e.g. Kay & Jost, 2003); and right-wing authoritarianism (RWA; Altemeyer, 1988).

I first looked at SDO. As discussed above, the SD-HARM (Dhont et al., 2016) model posits that SDO underpins both speciesism and human-human types of prejudice, and so I predicted that SDO would be correlated with speciesism and that SDO would account for the correlation between speciesism and human-human types of prejudice. SDO was measured using the SDO-6 scale (Pratto et al., 1994), which consists of 8 items including “Some groups of people

are simply inferior to other groups” ($\alpha = .93$). Second, I looked at political conservatism, which has been previously found to correlate with speciesism (Dhont et al., 2016). I measured conservatism using a standard measure (Graham, Haidt, & Nosek, 2009; Poteat & Mereish, 2012; Stern, West, Jost, & Rule, 2013) where participants indicated on two Likert scales from 1 (very liberal) to 7 (very conservative) the degree to which they identify as economically and socially conservatives respectively. The two measures were aggregated in the analysis.

Third, I looked at system-justification (Kay & Jost, 2003), which to my knowledge has not been investigated in the context of speciesism before. The scale consists of 8 items including “In general, you find society to be fair” ($\alpha = .86$). People who score high in system-justification tend to justify and defend the status quo, and given that the status quo places strict hierarchies amongst animals, I therefore assumed that they would also be more likely to defend the current speciesist norm. And fourth and finally, I looked at right-wing authoritarianism (RWA; Altemeyer, 1988), a construct tapping into the extent to which people adhere to established authorities and conventions and their hostility towards those who do not. The scale consists of 15 items including “The established authorities generally turn out to be right about things, while the radicals and protestors are usually just ‘loud mouths’ showing off their ignorance” ($\alpha = .92$).

Results showed that, supporting predictions, speciesism was significantly correlated with SDO ($r = .42, p < .001$), political conservatism ($r = .25, p < .001$), system-justification ($r = .25, p < .001$). The correlation between speciesism and RWA was positive, but not considered statistically significant under the adjusted Bonferroni alpha level of .006 ($r = .14, p = .03$). People that were more speciesist were also more likely to endorse hierarchies between group, report a more conservative political ideology, and more likely to engage in system justification. Recall that the SD-HARM model suggests that SDO is the common ideological root of both speciesism and human-human types of prejudice. In order to test this hypothesis, I conducted partial correlation analyses between speciesism and the other measures in which I controlled for SDO. And indeed,

when controlling for SDO I found that all partial correlations, but sexism and empathic concern became non-significant (*Table 2.3*).

Empathic concern and Actively Open-Minded Thinking. Finally, I addressed Aim 2c: identifying the relation between speciesism and empathic concern as well as actively open-minded thinking. I used the Empathic Concern scale, which forms part of the Interpersonal Reactivity Index (IRI; M. H. Davis, 1983) and consists of 7 items such as “When I see someone being taken advantage of, I feel kind of protective toward them” ($\alpha = .97$). As predicted the results revealed that speciesism correlated negatively with empathic concern ($r = -.46, p < .001$).

Next, I looked at actively open-minded thinking (Baron, 2000). The scale I relied on (Stanovich & West, 1997) consists of the subscales dogmatic thinking, categorical thinking, flexible thinking, counterfactual thinking, and openness. In total it consists of 40 items such as “A person should always consider new possibilities” ($\alpha = .92$). Speciesism correlated negatively with actively open-minded thinking ($r = -.17, p = .01$). This was mainly driven by dogmatic thinking ($r = .21, p < .001$) and flexible thinking ($r = -.16, p = .01$) and less so by openness ($r = .10, p = .07$), categorical thinking ($r = .10, p = .13$) or counterfactual thinking ($r = .03, p = .63$).

Vegetarianism. Fourteen of the 242 participants (6%) stated that they were vegetarian, and they were more likely to disagree with speciesism than the rest of the sample, $t(240) = 5.64, p < .001, d = 1.77$. Like in Study 2.1, speciesism correlated negatively with items capturing attitudes towards vegetarianism (second extracted factor of Study 2.1, see *Table 2.1*), $r(240) = -.31, p < .001$.

Demographics. There was a significant effect of gender such that men ($M = 3.82, SD = 1.30$) were more likely to show speciesist attitudes than women ($M = 2.91, SD = 1.26$), $t(239) = 5.46, p < .001, d = 0.71$. The negative correlation between speciesism and age did not reach significance, $r(240) = -.11, p = .08$, and there were no significant correlations between speciesism and level of education $r(247) = -.02, p = .80$, or income $r(247) = .10, p = .13$.

Table 2.3. Correlations between speciesism and other constructs as well as partial correlations (pr) when controlled for SDO (Study 2.3).

	r	pr
Social Dominance Orientation	.42***	
Racism	.32***	.01
Sexism	.41***	.20**
Homophobia	.17**	.04
Conservatism	.25***	.07
System Justification	.25***	.15†
Right-Wing Authoritarianism	.14*	-.08
Empathic Concern	-.46***	-.31***
Actively Open-Minded Thinking	-.17*	.05

Notes. Degrees of freedom were 243. † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$. Bonferroni adjusted alpha level was .006.

Discussion

As hypothesised, speciesism was positively associated with prejudicial attitudes such as racism, sexism, and homophobia. Speciesism was also positively associated with socio-ideological beliefs such as social dominance orientation, conservatism, system justification, and right-wing authoritarianism. Furthermore, speciesism was negatively associated with actively open-minded thinking and empathic concern. Note, however, that the correlations between speciesism and right-wing authoritarianism and actively open-minded thinking were relatively weak and not statistically significant under the adjusted Bonferroni alpha level of .006.

My findings lead me to conclude that speciesism can be considered a psychological prejudice analogous to other forms of prejudice. First, my results are consistent with Allport's

assumption (1954) of an underlying generalised prejudice: a person who exhibits one type of prejudice will likely also exhibit other types of prejudice, and importantly this extends to prejudice based on species-membership. This suggests that there is a common component of generalised prejudice that drives different types of specific prejudicial attitudes such as racism, sexism, homophobia as well as speciesism (Akrami et al., 2011). Second, my results support the SD-HARM theory (Dhont et al., 2016), which assumes that the same ideological roots that underpin human-human forms of prejudice also underpin speciesism. More specifically, the fact that SDO accounted for the relation between speciesism and other types of prejudice supports the notion that a general endorsement of social hierarchy and inequality drives these different manifestations of prejudice. The generalised prejudice and SD-HARM hypotheses are not mutually exclusive: research has shown, for example, that the generalised prejudice factor is strongly related to SDO, RWA as well as empathy (McFarland, 2010)—findings that are consistent with my results.

Further, my results suggest that people who endorse views on the political right are more likely to display speciesist attitudes. Considering that speciesism is the accepted social norm in Western society, it is not surprising that those who endorse the status quo are more likely to endorse speciesism. These findings are in line with previous research showing that people who consume meat score higher on SDO and RWA scales than vegetarians (Allen & Baines, 2002; Allen, Wilson, Ng, & Dunne, 2000).

The observation that speciesism correlated negatively with actively open-minded thinking supports my assumption that those who accept anti-speciesism are more willing to think beyond contemporary social norms. And as expected, speciesism correlated negatively with empathic concern. This correlation remained strong even when controlling for SDO, which suggests that empathic concern is an additional and independent route to reach anti-speciesist attitudes. This is in line with previous research showing that in addition to SDO and RWA low empathy is one of the roots of generalised prejudice (McFarland, 2010).

As in the first study, I found that male participants were more likely to endorse speciesist views than female participants. This confirms a similar finding that women have kinder attitudes towards animals than men (for a review see H. A. Herzog & Mathews, 1997; H. Herzog et al., 1991). Notably, previous research has pointed out that meat consumption is often related to masculinity. Rothgerber (2013) has found that for some men meat consumption makes them feel like “real men”. People who consume meat are perceived as more masculine (Ruby & Heine, 2011) and meat has become a metaphor for masculinity (Rozin, Hormes, Faith, & Wansink, 2012).

Study 2.4. Helping Via Donation Allocations

Thus far I have argued that speciesism is a psychological construct and developed a valid and reliable measure of speciesism (Study 2.1); confirmed that speciesist attitudes are consistent over time (Study 2.2); and shown that speciesism can usefully be considered a form of prejudice, being associated with other forms of prejudice and sharing common ideological roots with them (Study 2.3). Next, I turned to the more behavioural effects that speciesist attitudes might have by looking at whether speciesism can predict decision-making in the context of charitable giving. Previous work has exclusively focused on self-report items relating to feelings or beliefs about the relationship between humans and animals. In Studies 4 and 5 I addressed my third aim: to determine whether speciesist attitudes translate into observable decision-making and behaviour.

My initial assumption was that participants who scored high for speciesism would be willing to allocate more money to a charity that helps individuals of a “superior” species (rather than an “inferior” one). In particular, I hypothesised that speciesism would drive participants to a) help humans more than animals, b) help severely mentally disabled humans more than chimpanzees, and c) help dogs more than pigs. These examples capture three crucial aspects of speciesism (as defined in philosophy) in our society: a) the view that animals are morally inferior to humans; b) the view that (a) is true even when the humans and animals in question have similar

cognitive abilities; and c) the view that some animals are morally inferior to other animals (even when these animals have similar cognitive and emotional capabilities). As such, this study also serves as a philosophically grounded validation of my Speciesism Scale. The final items of the scale do not directly capture aspects (b) and (c). Yet, if it is true that there is a single psychological construct that captures speciesism in its entirety (as defined in philosophy), we should expect the scale to predict aspects (b) and (c) in addition to aspect (a).

With these parameters in mind, I hypothesised that speciesism would be able to predict donation allocations above and beyond established discrimination and pro-sociality measures such as SDO and empathic concern. One of the conclusions reached in my third study was that speciesism can be considered a sub-type of SDO as proposed by SD-HARM (Dhont et al., 2016). In this study I aim to show that speciesism despite its strong relation to SDO is able to explain discriminatory behaviour above and beyond it. I would expect speciesist behaviour to particularly be observed in contexts in which individuals of different species are contrasted. In such contexts, we would expect speciesism to give more specific predictions than the more generalised concept of SDO.

Finally, I was also interested in the link between speciesism and people's explicit beliefs about the intelligence and sentience of animals. It is possible, for example, that people are not treating certain animals differently because of speciesism, but simply because they believe that they are less intelligent or less able to suffer. I hypothesised that explicit beliefs about intelligence and sentience were likely to be associated with speciesist views but would not be able to fully explain the differential treatment between different species. Instead, I expected that speciesism itself—moral discrimination based on species-membership alone—would be able to explain differences in donation allocations above and beyond explicit beliefs about intelligence and sentience.

Method

Participants. 140 US American participants took part in the study online via MTurk (110 female; $M_{age} = 33.60$, $SD = 10.54$), and received \$1.70 for their participation. All participants successfully passed the two attention check questions, so nobody was excluded. I deemed a sample size of 140 participants appropriate, as my a priori power analysis showed that 135 participants were required to detect an anticipated medium effect size of $f = 0.15$, taking an alpha of .05 and power of .80 (Cohen, West, & Aiken, 2014).

Materials and Procedure. Participants were presented with three independent scenarios of donation allocation questions in a random order. In each scenario, one of the charities focused on helping a “superior” type of species (humans, mentally severely disabled humans, or dogs), while the other focused on helping an “inferior” type of species (animals, chimpanzees, or pigs).

The two charities were presented as follows in the first donation scenario:

Charity A: This charity focuses on helping *animals*. On average, \$10 will keep one individual from pain and suffering for one day.

Charity B: This charity focuses on helping *humans*. On average, \$10 will keep one individual from pain and suffering for one day.

In the second scenario the words “animals” and “humans” were replaced by “chimpanzees” (Charity A) and “mentally severely disabled humans” (Charity B), in the third scenario by “pigs” (Charity A) and “dogs” (Charity B) respectively. In all three scenarios, participants were asked to allocate 100 donation points, which were independent of participants’ payment, between the two charities. They were told that their decisions would have a real-world effect, as the experimenters would donate \$100 to the charities in proportion the responses received from all participants (which I did).

After participants responded to the donation allocation questions, they were presented with the Speciesism Scale and the following additional measures: Social Dominance Orientation

was included because of the strong link found to speciesism in the previous study. Attitudes Toward Disabled Persons (Yuker, Block, & Young, 1970; from now on referred to as ableism) was included in order to control for discrimination against disabled people in the donation case involving mentally severely disabled humans. Empathic concern (IRI; M. H. Davis, 1983) was included because it showed a strong negative correlation with speciesism in the previous study. Perspective taking, Other-oriented moral reasoning and Mutual concerns moral reasoning scales of the pro-social personality battery (Penner, Fritzsche, Craiger, & Freifeld, 1995) were included because I assumed that these pro-sociality traits could potentially play a role in reducing prejudice.

Further, participants were presented with a list of six different types of individuals: humans in general, animals in general, chimpanzees, mentally severely disabled humans, dogs, and pigs. First, participants were asked to indicate to which extent each type of individual was capable of experiencing physical pain and fear (aggregated and from now on referred to as ‘suffering capability’). Then, participants were asked to indicate how intelligent they perceived each type of individual to be. By framing this task as a biology quiz, participants were told that they should be as accurate as possible in their assessment of these traits. Finally, participants were asked to respond to demographic items.

Results

People allocated significantly more to the “superior” than to the “inferior” species charities in all three cases (*Figure 2.1*). As the allocations to the respective charities were negatively proportionate to each other, one sample t-tests were conducted to compare whether allocations significantly differed from an equal distribution of 50 donation points each. Participants allocated more to help humans ($M = 67.92$, $SD = 26.36$) than animals ($M = 32.08$, $SD = 26.36$), $t(139) = -8.04$, $p < .001$, $d = .67$, more to help mentally severely disabled humans ($M = 72.46$, $SD = 23.61$) than to chimpanzees ($M = 27.54$, $SD = 23.61$), $t(139) = 11.26$, $p < .001$, $d = .95$, and more to help

dogs ($M = 68.57, SD = 31.43$) than pigs ($M = 31.43, SD = 19.43$), $t(139) = -11.29, p < .001, d = .59$.

Speciesism correlated positively with SDO ($r = .35, p < .001$) and negatively with empathic concern ($r = -.26, p < .001$) and other-oriented ($r = -.23, p < .05$) and mutual concern moral reasoning ($r = -.20, p < .05$). However, it did not correlate significantly with ableism ($r = -.04, p = .60$) and perspective taking ($r = -.09, p = .32$). Speciesism correlated negatively with most measures of beliefs about suffering capability and intelligence of animals (but not humans), and strongest for members of “inferior” species (Table 2.4). Participants believed that humans were more intelligent ($t(139) = 8.25, p < .001$) and more capable of suffering ($t(139) = 26.63, p < .001$) than animals. However, despite allocating more money to help mentally severely disabled humans than chimpanzees, they believed the former were less intelligent ($t(139) = 4.98, p < .001$) and—although only on a descriptive level—less capable of suffering ($t(139) = 1.71, p = .09$) than the latter. Also, participants believed that dogs were more intelligent ($t(139) = 2.22, p = .03$) and more capable of suffering ($t(139) = 5.58, p < .001$) than pigs.

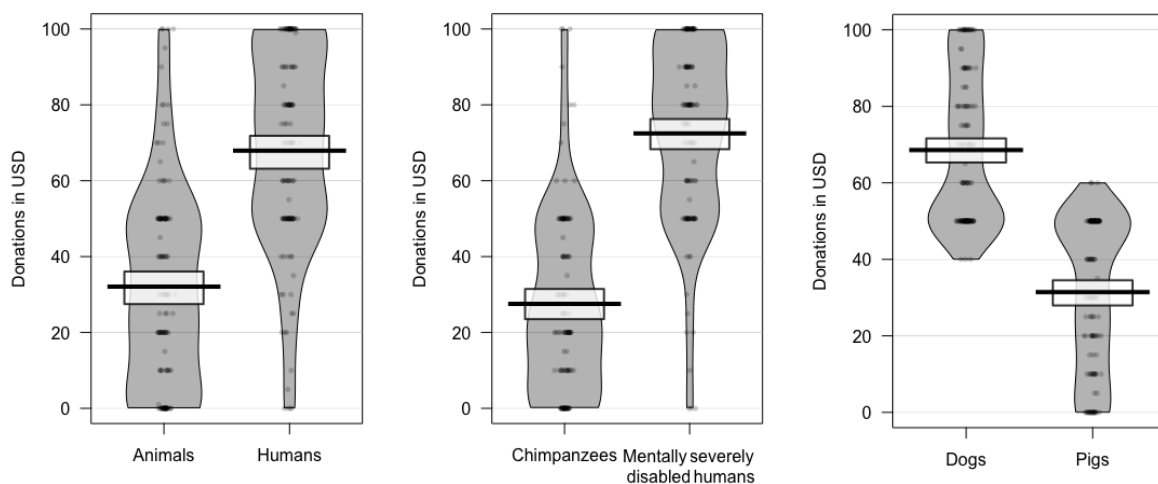


Figure 2.1 Donation allocations to the charities (Study 2.4). In all three scenarios allocations significantly differed from an equal distribution of 50 donation points each. Black points represent

raw data, horizontal bars represent means, rectangles represent confidence intervals, and “beans” represent smoothed densities.

Table 2.4. Beliefs about intelligence and suffering capability (ratings on a scale from 0 to 100) and correlations of these beliefs with speciesism (Study 2.4).

	Beliefs in <i>M (SD)</i>		Correlation with speciesism	
	Intelligence	Suffering capability	Intelligence	Suffering capability
Humans	87.79 (12.88)	96.88 (6.82)	-.04	-.15
Animals	39.71 (19.71)	84.20 (20.39)	-.27***	-.29***
Mentally severely disabled humans	54.14 (26.44)	89.96 (15.31)	0.6	-.10
Chimpanzees	66.84 (22.61)	91.92 (13.37)	-.20*	-.25**
Dogs	55.09 (21.25)	91.33 (12.34)	-.18*	-.28***
Pigs	51.43 (25.70)	87.07 (18.38)	-.23**	-.29***

Notes. * $p < .05$. ** $p < .01$. *** $p < .001$.

Hierarchical regression analyses were conducted in order to assess whether speciesism is able to uniquely predict participants’ allocations. Demographics (gender, age, education, income) were entered at Step 1; Social psychological discrimination and pro-sociality constructs (SDO, empathic concern, perspective taking, other-oriented moral reasoning, mutual concerns moral reasoning, and ableism in the case involving mentally severely disabled humans) were entered at Step 2; Beliefs about suffering capability and intelligence of the respective two entities per case were entered at Step 3; Speciesism was entered at Step 4.

Tests of collinearity indicated that multicollinearity was not a concern (all variance inflation factors (VIF) were below 2.7). The analyses revealed that speciesism was able to predict the amount of allocations in all three cases significantly above and beyond demographic, discrimination and pro-sociality measures and perceived suffering capability and intelligence. Thus, despite the fact

that people with high scores on speciesism believed animals' suffering capability and intelligence were reduced, speciesism significantly predicted allocations above and beyond these beliefs (*Table 2.5*)

Table 2.5. Hierarchical regression predicting donation allocations to “superior” individuals (humans, mentally severely disabled humans, dogs) vs. “inferior” individuals (animals, chimpanzees, pigs), displaying the standardised coefficients (beta) (Study 2.4).

Predictor measures	Humans vs. Animals				Mentally Severely Disabled Humans vs. Chimpanzees				Dogs vs. Pigs			
	Step 1	Step 2	Step 3	Step 4	Step 1	Step 2	Step 3	Step 4	Step 1	Step 2	Step 3	Step 4
ΔR^2	.02	.04	.11**	.17***	.03	.04	.17***	.07***	.01	.05	.04	.07**
Gender	-.01	-.03	-.14	-.15†	.06	.02	.16	.17†	.05	.03	.03	.01
Age	-.10	-.13	-.09	-.03	.12	.13	.08	.04	-.07	-.09	-.07	-.01
Education	.03	.03	.03	.08	.12	.12	.12	.09	.02	.03	.00	.04
Income	.05	.06	.04	.00	-.06	-.08	-.13	-.09	-.07	-.07	-.07	-.10
SDO		-.08	.04	.06		.02	.03	-.08		-.15	-.13	-.01
Empathic concern		.07	-.11	-.06		.00	.05	.12		.14	.06	-.01
Perspective taking		-.14	.04	-.01		.19	.21†	.14		.12	.15	.22†
Other-oriented		-.08	-.11	-.13		.03	.04	.08		-.11	-.13	-.18
Mutual concern		.23‡	-.07	.21†		-.08	-.14	-.14		-.11	-.06	-.07
Ableism						.12	.04	.05				

Suffering “superior”	.05	.04	.18†	.17†	-.43*	-.43*
Suffering “inferior”	.14	.06	-.21*	-.18†	.47*	.42*
Intelligence “superior”	-.26*	-.21*	.29**	.25**	.02	-.01
Intelligence “inferior”	.25**	.13	-.27**	-.22*	.05	.01
Speciesism		-.46***		.30***		-.27***

Notes. † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$.

Discussion

In Study 2.4, I looked at whether speciesist attitudes translate into decisions about donation allocations. On average, participants allocated higher donations to a) a charity focusing on helping humans than to a charity focusing on helping animals; b) a charity focusing on helping mentally severely disabled humans than to a charity focusing on helping chimpanzees; and c) a charity focusing on helping dogs than to a charity focusing on helping pigs. Speciesism was able to explain these allocation decisions above and beyond demographic measures; discrimination and pro-sociality measures (including SDO and empathic concern); and measures of explicit beliefs about suffering capability and intelligence. I conclude that speciesism makes a unique contribution to our understanding of how people judge and treat individuals of different species over established social and psychological constructs. It uniquely captures an important but neglected aspect of discrimination and pro-sociality. These findings further demonstrate the (philosophical) content validity of the Speciesism scale. The psychological construct that the Speciesism Scale measures reliably predicts all major aspects of speciesism in our society as defined in philosophy.

People who displayed speciesist attitudes were more likely to believe that animals have a reduced capability to suffer (particularly animals of “inferior” species) as well as have lower intelligence. This is consistent with studies that have found people tend to ‘de-mentalize’ animals classified as food (Bastian et al., 2012) and that the extent of moral value they attribute to animals is related to the animal’s degree of intelligence and sentience (Piazza & Loughnan, 2016). My study, however, did not investigate the causal direction of this association. It could be that speciesism drives people to attribute lower mental capabilities to animals or that lower attributed mental capabilities drive speciesism. In the present study, people’s explicit beliefs of the ability to suffer and the level of intelligence did explain their donation allocations to some extent. However, speciesism was able to explain it above and beyond these beliefs in all three cases. Interestingly, my study showed that in some cases people morally value individuals of certain species less than

others even when knowing that the former (e.g. chimpanzees) are more intelligent and more sentient than the latter (e.g. mentally severely disabled humans). These findings confirm my hypothesis that the differential treatment of animals is not motivated by people's explicit beliefs about sentience and perceived intelligence alone; it is at least in parts motivated by speciesism.

It is important to note that participants in my study were explicitly asked to rate the suffering capability and level of intelligence as biologically accurate as possible. In real life settings, in contrast, motivated reasoning can distort these beliefs—as demonstrated, for example, by studies investigating the meat paradox (Bastian & Loughnan, 2017). People's perceptions of these mental properties are thus very flexible and depend on the context and motivational incentives. Therefore, even though my study demonstrates that differential treatment of individuals of different species cannot be fully explained by people's explicit and accuracy-motivated beliefs about their mental capabilities, it is still possible that people's intuitive and biased perceptions significantly drive this tendency.

A possible criticism of the study is that even though participants were aware of the fact that their decisions were not just merely hypothetical, they did not have to invest personal resources in order to help the respective individuals. Instead, they decided how to allocate the available resources of the experimenter. In this light, their responses are better described as resource allocation or decision-making, but not necessarily as truly behavioural. In the next study, I aimed to address this issue by relying on a response technique, which required participants to invest their personal resources.

Study 2.5. Helping Via Time Investment and Food Choice

In Study 2.4, I demonstrated that speciesism predicts decision-making about donation allocations above and beyond other related psychological constructs. I conducted Study 2.5 with three purposes in mind. First and foremost, I aimed to replicate the finding that speciesism is able

to predict helping. Second, instead of decisions regarding allocations of donations, I instead asked participants in this study whether they would invest their personal time to help individuals, and I explored whether speciesism predicts food choices. In doing so I introduce measures that are clearly behavioural. Third, I aimed to confirm that speciesism was able to predict real-world behaviour in an offline, university context relying on a student sample.

In this study, I asked participants to read promotional material from a charity and asked for their feedback on how it could be improved. Each participant had to choose from two charities; one focused on helping animals and the other on helping humans. I hypothesised that participants scoring low on speciesism would be more likely to invest time in the animal charity. I further hypothesised that speciesism would be able to predict which charity participants helped above and beyond social dominance orientation, empathic concern, perceived suffering capability and intelligence of the type of individuals, and demographic data.

I also included an exploratory test related to food choices. At the end of the study, participants could choose either a meat snack or a vegetarian snack as an additional reward for participating in the study. Based on my findings from the previous studies showing that speciesism correlates negatively with ethical vegetarianism I hypothesised that those participants scoring high in speciesism would be more likely to choose the meat snack. This exploratory test can be seen as a proxy for meat consumption, but it should be interpreted with caution as I measured food choice at a certain point in time only, not general meat consumption.

Method

Participants. 102 students from a British university took part in the study and received a payment of £5 for their participation. Three participants were excluded for failing an attention check, leaving a final sample of 99 students (55 female; $M_{age} = 25.52$, $SD = 4.88$). This sample size meets the recommendation by Long (1997), who suggests that a minimum sample of 100 should be considered for logistic regression, and my sample size is comparable to work by Crimston et al.

(2016), who conducted a structurally very similar study to ours and recruited the same number of participants as I did.

Materials and Procedure. Participants completed the study in “paper & pencil” form. They first read a general information sheet and signed a consent form. They then completed the Speciesism Scale, SDO scale, and Empathic Concern scale. Next, they were asked to indicate their beliefs in the intelligence and suffering capability of humans, chimpanzees and dogs. And finally, participants responded to demographic questions (age, gender, and study subject).

My method for measuring participants’ willingness to invest time to help a certain cause was inspired by Crimston et al. (2016). On a further sheet, participants were given instructions about a seemingly unrelated task. Participants were informed that they should provide feedback to one of two non-profit organizations (from now on referred to as charity) on how to improve their campaigns. They were told that both charities were specifically seeking feedback from students.

The two charities were described as follows:

The Nonhuman Rights project: The Nonhuman Rights Project, an organization working towards achieving actual legal rights for chimpanzees.

Homeless Shelters: Homeless Shelters, an organization providing temporary residence for people who can no longer afford to pay rent themselves.

Participants were presented with two sealed envelopes labelled with the respective names of the charities. Promotional material from the charities was contained inside the envelope. Participants were instructed to open one envelope only, and to provide written feedback about the campaign. Participants were told that their feedback would be forwarded to the charity of their choice.

At the end of the study, after receiving the payment, participants were offered a snack as an additional reward. They were presented with two meat options (beef jerky, pork crackling) and

two vegetarian options that clearly did not include any animal products (salted peanuts, rice crackers) and were asked to choose one.

Results

As predicted, participants who chose to invest their time to help the animal charity (Non-Human Rights Project) scored significantly lower on speciesism ($N = 33$, $M = 2.73$, $SD = .96$) than those who chose the human charity (Homeless Shelters) ($N = 66$, $M = 3.31$, $SD = .83$), $t(97) = -3.16$, $p = .002$. However, participants who chose to help Homeless Shelters scored higher on empathic concern ($M = 5.58$, $SD = .75$) and lower on SDO ($M = 2.14$, $SD = 1.04$) than participants who chose the animal charity (empathic concern: $M = 5.18$, $SD = .89$ $t(97) = -2.32$, $p = .02$, SDO: $M = 2.72$, $SD = .89$ $t(97) = 2.76$; $p = .01$). Notably, SDO was left skewed and therefore not normally distributed.

Hierarchical logistic regression analyses were conducted in order to assess whether speciesism was able to predict participants' choice of the charity above and beyond the other measures. Demographics (gender, age, study subject) were entered at Step 1; SDO and empathic concern were entered at Step 2; Beliefs about suffering capability and intelligence of chimpanzees and humans were entered at Step 3; Speciesism was entered at Step 4. The analyses revealed that speciesism was able to predict participants' time-investment choices significantly above and beyond all other factors listed. Results are shown in *Table 2.6*.

Table 2.6. Hierarchical logistic regression predicting whether participants chose the human or animal charity), displaying Exp(B) (Study 2.5).

Predictor measures	Step 1	Step 2	Step 3	Step 4
R ₂	.00	.09*	.11	.23***
Gender	1.04	.85	.99	1.23
Age	1.02	1.01	1.01	1.02
SDO		.64†	.64†	.62†
Empathic concern		1.58	1.69†	1.97†
Suffering humans			1.00	1.01
Suffering chimpanzees			.996	1.00
Intelligence humans			1.02	1.00
Intelligence chimpanzees			.98	.99
Speciesism				3.43**

Notes. † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$.

Further, participants who chose the meat snack ($N = 26$, $M = 3.47$, $SD = .90$) scored also significantly higher on speciesism than participants who chose the vegetarian snack ($N = 72$, $M = 2.98$, $SD = .88$), $t(96) = -2.40$, $p = .018$. Hierarchical logistic regression analysis predicting the snack participants chose analogously to the above analysis was close to being significant, $R_2 = .23$, $p = .071$.

Discussion

This study demonstrated that speciesism is able to predict how people will invest their time when faced with a choice between helping animals or humans. After controlling for demographics, empathic concern, SDO, and beliefs about suffering capability and intelligence of humans and

chimpanzees, each one-point increase in speciesism increased the likelihood of people choosing to help the human charity by 195%. Significant results supporting this finding were also found in an online pilot study (N = 51) I conducted prior to the student sample study.

Interestingly, people who chose the animal charity showed less empathic concern than people who chose the human charity. Empathic concern, therefore, seems to be a stronger predictor for helpful behaviour towards humans than towards animals (when one is only faced with these two options only). This is in spite of the negative correlation between speciesism and empathic concern. One possible explanation for this result is that people feel empathic concern to those who are close to them, and people feel more closely connected to other humans than to animals. Another explanation might be that fighting homelessness triggers more empathy than establishing animal rights because the former is perceived as more imminent and more directly tractable than the latter.

Other results were also puzzling at first. For example, people who chose to help the animal charity also scored higher in SDO than people who chose to help the human charity. At first glance, this appears inconsistent with my previous findings. However, we should be cautious to attribute much weight to this result, as SDO was strongly left skewed in the university student sample. In contrast, in the online pilot study SDO and empathic concern were normally distributed and no correlations between the charity chosen and SDO and empathic concern were revealed.

This study further demonstrated that speciesism can predict real-world food choices. The lower people scored on speciesism, the more likely they were to choose the vegetarian snack. This finding confirms the result of the online pilot study in which participants were asked a hypothetical question about what food they would choose. The food-choice element of the study was exploratory, and there were many other factors that could have affected participants' choice of snack. These might include the social demand effect, different tastes, and the time of day, among others. A further limitation of this study was that I did not control for vegetarianism. Due to the

relation between speciesism and vegetarianism, it is not clear to which extent food choices were driven by speciesism versus vegetarianism. However, given that only 6% of US participants of Study 2.1 indicated that they were vegetarian, and recent polls suggest only 2% of the British public is consistently vegetarian (Bates et al., 2014), only little variance in food choices could have been explained by vegetarianism alone. Given that 74% of my participants chose the vegetarian snack, yet estimates suggest no more than 6% of the population to be vegetarian, it seems clear that many non-vegetarians must have chosen the vegetarian snack and there are many reasons for why this might be—lower scores in speciesism is likely to be one of them. Future research can hopefully shed more light on the relation of speciesism, vegetarianism, carnism, and actual food choices.

A general methodological limitation of the study is the possibility that demand effects or consistency motives could have driven the correlation between responses to the scale and responses to the tasks. Participants might have guessed what the aim of the study was after responding to the scale and might have been motivated to remain consistent in the next task. Future studies could attempt to alleviate this problem by trying to obfuscate the purpose of the study, introducing time delays between the tasks, or by other means.

General Discussion

Throughout its history, speciesism has been a concept largely confined to philosophy. Ultimately, however, speciesism is a hypothesis about human psychology; it is the assumption that people morally discriminate against individuals on the basis of their species-membership. The purpose of this chapter was to transition speciesism from a philosophical hypothesis to a matter of empirical psychological investigation. These investigations have only just begun in recent years and with this chapter I make the following three contributions to that transition.

First, I presented the *Speciesism Scale*: A theoretically driven and empirically validated explicit measure of speciesism with high internal consistency and test-retest reliability. This scale confirms

that speciesism is an accurately measurable, stable form of prejudice with high interpersonal differences.

Second, I demonstrated that, as originally proposed by philosophers, speciesism can be considered a form of prejudice. I have found that speciesism is psychologically related to human-human types of prejudice such as racism, sexism, and homophobia. This is consistent with the generalised prejudice theory, which assumes that prejudice towards various targets is significantly correlated and can be explained by an underlying generalised prejudice factor (Akrami, Ekehammer, & Bergh, 2011). My finding that speciesism is strongly related to SDO confirms SD-HARM (Dhont et al., 2014), which assumes that speciesism and human-human types of prejudice are underpinned by the same socio-ideological beliefs, namely a general endorsement of social hierarchy and inequality. The fact that speciesism shares these psychological properties with other phenomena referred to as prejudice seems to strengthen the case that speciesism can be referred to as prejudice as well. I note that in comparison to human-human forms of prejudice, speciesism is the dominant and explicitly accepted social norm and ideology in current Western societies. Consequently, people who endorse the current status quo and people on the political right tend to score high on speciesism, whereas actively open-minded thinking seems to facilitate questioning that norm. I also identified lower levels of empathic concern as an additional factor associated with speciesism.

Third, I demonstrated that speciesism can predict behaviour above and beyond existing constructs of prosociality as well as perceived mental capacities. In doing so, I believe I am the first to systematically show effects of speciesist attitudes on actual, observable behaviour. In two studies I found that speciesism predicts people's willingness to help humans and "superior" animals such as dogs (rather than "inferior" animals such as pigs), in terms of allocating donation money and investing time. I also found that speciesism predicts people's (meat vs. vegetarian) food

choices. This indicates that speciesism captures an important but neglected aspect of both discrimination and prosociality.

Limitations and future research

One might argue that the scale does not clearly distinguish between speciesism and the broader phenomenon of moral anthropocentrism, i.e. the view that humans are morally more valuable than animals for *any* reason, such as pure species-membership or more advanced mental capacities. Indeed, some of the items (but not all) do not completely rule out the possibility—at least in principle—of also capturing a general tendency of devaluing mentally less advanced beings. For example, it is possible that people who consider it morally acceptable to trade animals like possessions do so not purely because of species-membership but also because they believe that it is in general permissible to trade beings that only have basic mental capacities. However, this seems unlikely because almost all people consider it wrong to trade humans with similarly basic mental capacities, which means that the differential treatment must be related with species-membership itself. The same applies to all other items, such as the item about keeping animals in circuses for entertainment. In fact, all items were devised in a way such that people would likely disagree with them if they involved mentally basic humans instead of animals. Therefore, while not all items completely ruled out the possibility of also capturing a general tendency of devaluing mentally less advanced beings, I expect that the Speciesism Scale almost exclusively captures speciesism. In Chapter 3, I will investigate the relation between speciesism, perceived mental capacities, and the broader phenomenon of moral anthropocentrism in more detail.

My investigation of speciesism focused on the general tendency to morally discriminate between individuals on the basis of their species-membership. However, I did not explicitly analyse how this occurs. It is possible that—in line with research on stereotypes and social categorizations (Haslam, 1997)—people place individuals of certain species into different categories such as food, pets, entertainment, equipment, pest, wild animals, humans, etc. These culturally and historically

contingent categories might define the general connotation, perceptions, values, and common properties associated with their individuals. For example, animals in the food category (e.g. cows) seem to be associated with low moral value as well as perceptions of low mental capabilities. Animals of the pest category (e.g. rats) seem furthermore to be associated with disgust-eliciting characteristics such as being dirty. The categories ultimately define (or are defined by) how we relate to the respective individuals. And while usually whole species are placed into certain categories (i.e. speciesism), exceptions can be made; for instance, one's personal pet pig will be placed into a different category than all other pigs. Future research will hopefully further investigate the specifics of the different kinds of categories that speciesism seems to rely on.

Similarly, more research is needed to understand the cognitive mechanisms that facilitate speciesist attitudes and practices. In this regard, it is likely that the psychological process of moral disengagement (Bandura, 1999) plays a critical role. Bandura's theory of moral disengagement proposes that cognitive restructuring mechanisms allow us to view inhumane practices as benign by disabling moral self-condemnation. These mechanisms may include moral justification, euphemistic language, displacement of responsibility, dehumanization, and a disregard for the victims' suffering. Moral disengagement theory has not yet been applied to speciesism, but it is a plausible explanation for tolerating practices such as factory farming (Graça, Calheiros, & Oliveira, 2016).

Moral philosophers have reasoned that anti-speciesists are morally compelled to adopt ethical vegetarianism (Singer, 1975, 1979). Descriptively, however, I have found that empirically this is not always the case. People do not consistently bridge the gap between caring about animals and consuming animal products (i.e. meat paradox). My results from Study 2.1 and 2.3 suggest that speciesism and ethical vegetarianism are psychologically distinct to some extent. Overall, however, the two are strongly linked and indeed as Study 2.5 suggests speciesism predicts whether people choose to eat meat or a vegetarian option. It has been suggested that a set of ideological beliefs,

referred to as *carnism*, legitimize meat consumption in our society (Joy, 2011; Monteiro et al., 2017). Future studies in this area will now be able to use the Speciesism Scale to investigate the relations and causal mechanisms between speciesism, carnism and the meat paradox with greater rigor and accuracy.

Further research is needed on how speciesism varies between cultures. Amiot and Bastian (2015) have summarised existing studies that may help to point the way. Their work focused on cross-cultural attitudes towards animals. They found, for example, that attitudes towards pet keeping vary strongly between cultures. In order for the Speciesism Scale to facilitate cross-cultural research, however, it might need to be adapted for cultures where certain animals are categorised differently to Western norms.

I hope that the Speciesism Scale will prove to be a useful instrument for researchers to measure speciesism in different populations and to detect correlations between speciesism and other constructs. In the next two chapters, I will build on the findings of this chapter: first, by trying to understand more deeply what psychological factors underpin speciesist tendencies (Chapter 3); second, by investigating how speciesism translates into deontological constraints against harm (Chapter 4).

3. A Multi-Factorial Framework for Moral Anthropocentrism

In Chapter 2, I showed that speciesism is a stable and measurable psychological construct: people value humans more than animals due to species-membership. In this chapter, I give a deeper explanation of speciesism and disentangle its different sub-factors. At the same time, I zoom out and ask whether there are other factors, next to speciesism, that could explain the more general phenomenon of moral anthropocentrism. Relying on work in moral philosophy, I discuss possible grounds that humans could have for valuing humans more than animals. I then experimentally test which factors could explain why people value humans more than animals. My primary aim is to provide a framework for moral anthropocentrism in order for future research to investigate the psychology of its sub-factors in greater detail. While Chapter 2 relied on a conventional social psychological approach, in this chapter I rely on an approach that partly resembles experimental philosophy. In particular, I make use of hypothetical thought experiments that are required to disentangle the fine-grained facets of moral anthropocentrism.

There are two main ways in which moral anthropocentrism might be explained. The first is that people value humans more than animals because humans possess (or typically possess) morally valuable properties that other animals lack but which, potentially, other beings could have. A salient example is the view that humans are more valuable than animals because they have *advanced mental capacities* which animals lack. These include complex language and thoughts, the ability to plan for the future, to make deliberate and autonomous judgments (including moral evaluations), and complex emotions. Capacities such as these are often regarded as being required to be considered a ‘person’ that is granted full moral status (e.g. Kant, 1785; Quinn, 1984). On this

hypothesis—which I call the *Mental Capacity View*—people value humans more than animals because only humans have advanced mental capacities.

The *Mental Capacity View* comes in two versions. The most straightforward version is what I call the *Individual Mental Capacity View*, which says that individual humans are viewed as morally more important than individual animals because humans have more advanced mental capacities. If this view is the primary source of moral anthropocentrism, we would expect people to regard *non-human* individuals that have mental capacities similar to those of humans to be as morally important as humans. One problem for that hypothesis, however, is that some humans are treated as possessing full moral status despite having less advanced mental capacities than many animals. Familiar examples are infants and severely cognitively impaired people—often referred to as ‘marginal cases’. One possible reason why such people are granted full moral status is that they belong to a species whose members *typically* have advanced mental capacities. This more complex hypothesis—which I call the *Species-Typical Mental Capacity View*—says that people value humans more than animals because humans *typically* have more advanced mental capacities than animals.

The second class of hypotheses says that people value humans more than animals not because of humans possess certain valuable properties that other animals lack (but which other beings could potentially have) but directly because humans are members of the species *Homo Sapiens* (i.e. speciesism; cf. Singer, 1975, 1979). As I have shown in Chapter 2, speciesists assign higher or lower moral status to individuals purely on the basis of their species-membership. Discussions of speciesism, however, do not usually distinguish two very different forms that speciesism can take. The first is that people accept what I call *Species-Relativism*: the view that all individuals should prioritise members of their own species over others because we have special duties to members of our own species (Pugh, Kahane, & Savulescu, 2013; B. Williams, 2009). Notice that while this general ‘relativist’ principle says that humans should prioritise humans over non-humans, it *also* says that members of other species should prioritise *their* fellow species-

members—potentially even at the expense of humans. Hence on this view, how someone should prioritise between different species depends on the species they belong to.

I call the second form of speciesism *Pro-Human Bias*: the view that humans are more valuable than animals simply because they are humans, i.e. an *absolute* moral claim. Unlike Species-Relativism, Pro-Human Bias says that even non-humans should prioritise humans over members of their own species, even when the humans have the same, or indeed inferior, mental capacities. Such a distinctive bias in favour of humans *qua* humans could be called a prejudice.

Clarifying which, if any, of these factors underlies moral anthropocentrism could bear on the ethical evaluation of moral anthropocentrism, whether by showing that dominant philosophical defences of it are disconnected from what actually drives this view in most people, or by revealing that moral anthropocentrism is driven by factors—such as Pro-Human Bias—which cannot be given a plausible moral justification, and which can be described, without qualification, as a prejudice.

The Present Research

Discussion of what I call moral anthropocentrism often assumes that it has a single source. The existing evidence discussed in Chapter 1, however, suggests that multiple factors might be at work. My guiding hypothesis was therefore that both the Mental Capacity View and speciesism play a role. However, to disentangle these different factors, we cannot just look at cases contrasting humans with existing animals, since all existing animals have less advanced mental capacities than the typical human. Instead, we also need to look at more unusual scenarios involving mentally advanced animals or extraterrestrials. While such scenarios are purely hypothetical, they are a necessary tool to uncover the factors that underlie moral anthropocentrism. They have the further advantage that, although such imagined intelligent beings are familiar from popular culture, there is no prior conventional view about their moral status that might shape the responses of participants.

All my studies involve moral dilemmas where participants are asked to prioritise between two individuals that both require urgent help. In Study 3.1, I test the Individual Mental Capacity View by presenting participants with moral dilemmas involving humans of different mental capacity levels and monkeys (including monkeys with advanced mental capacities). In Study 3.2, I test the Species-Typical Mental Capacity View according to which moral status is determined by the typical mental capacity level of a species. Here I rely on moral dilemmas that contrast humans with members of hypothetical species (extraterrestrials) with either typically advanced or typically basic mental capacities. In Study 3.3a and 3.3b, I test the Species-Relativism and the Pro-Human Bias hypotheses, by asking who humans and non-humans, respectively, should prioritise in identical moral dilemmas that involve humans and non-humans. In Study 3.4, I extend my question to contexts involving harm, and contrast attitudes to animals vs. to other non-humans.

Open Science. Reports of all measures, manipulations, and exclusions, as well as all data, analysis code, and experimental materials are available for download at: https://osf.io/6ncb8/?view_only=dc04af067cc14c4cada6a6b3ab1703a8. Studies 3.2 and 8b were pre-registered on the Open Science Framework (see links below).

Ethics Statement. For all studies, relevant ethical guidelines were followed, and the research was approved through University of Oxford's Central University Research Ethics Committee, with the reference number MS-IDREC-R56657/RE001.

Study 3.1. Individual Mental Capacity Level

In Study 3.1, I aimed to test whether people value humans more than animals primarily because humans have advanced mental capacities that allow for complex thoughts and emotions, deliberate planning, and autonomous choice, whereas animals do not. More generally, I wanted to examine how differences in level of mental capacity affect moral status.

If humans are judged to be more morally valuable than animals because they possess more advanced mental capacities, then hypothetical animals whose mental capacities are at the human level should be judged to be as morally valuable as humans. However, I hypothesised that participants would disagree. Instead, I hypothesised that people would continue to value humans more than animals *even* if the animals were said to have human-level mental capacities. One reason for this hypothesis is that people seem to value humans more than animals in cases where both humans and animals have similarly *basic* mental capacities. For example, people are more willing to help a severely cognitively impaired human than a chimpanzee that has the same mental capacities (Chapter 2).

Another implication of the Individual Mental Capacity View is that animals with *above* human-level mental capacities should be viewed as morally *more* valuable than humans. However, my hypothesis was again that people would continue to value humans more than animals even if the animal has above-human-level mental capacities.

Together, these two predictions meant that I did not think that the Individual Mental Capacity View could fully explain moral anthropocentrism. However, I also predicted that everything else equal, people would find more mentally advanced individuals more morally valuable. Hence, I hypothesised that the Individual Mental Capacity View was nevertheless still part of the explanation of moral anthropocentrism.

Methods

Participants. I recruited 152 US American participants online via MTurk. They received \$0.85 in payment (in line with US minimum wage) for their participation. Forty-eight were excluded for failing at least one comprehension check, leaving a final sample of 104 people (59 female; $M_{age} = 35.55$, $SD_{age} = 10.70$). There were three comprehension checks that asked about the definitions of the three levels of mental capacity that the monkeys had. The fact that such a large fraction of participants failed at least one comprehension check I can only explain by the counter-

intuitiveness and abstractness of the descriptions. Since this was the first study, I did not know what effect sizes to expect. A power analysis showed that a sample of 90 participants or more would be adequate in order to detect small to medium effect sizes of $d = .3$ in one-sample t-tests, with an alpha of 0.05 and a power of 0.8.

Materials and Procedure. Employing a within-subject design, all participants were presented with the same 15 different dilemmas in randomised order. All dilemmas featured two individuals (the potential beneficiaries) that were pitted against each other. Participants were told that both individuals would experience suffering and that they had to decide which one they would prioritise helping if only one could be helped. The two individuals were either a human beneficiary and a monkey beneficiary (nine dilemmas), two human beneficiaries (three dilemmas) or two monkey beneficiaries (three dilemmas). The individuals either had basic (animal-level), advanced (human-level), or extremely advanced (above-human-level) mental capacities⁴. Consequently, some dilemmas involved fantastic cases such as monkeys with human-level or even above-human-level mental capacities. In each dilemma, participants had to indicate which individual they would prioritise helping (by giving a painkiller medicine to relieve them of pain) on a 7-point scale (1 = *Definitely A*, 4 = *Flip a coin to decide*, 7 = *Definitely B*), where A and B refer to the two individuals respectively and differ in each dilemma. Finally, participants responded to the Speciesism Scale (Chapter 2), the Oxford Utilitarianism Scale (OUS; Kahane et al., 2018), and demographic questions. The OUS is a scale measuring individual differences in utilitarian tendencies. It consists of two largely uncorrelated subcomponents of utilitarianism. While the former captures permissive attitude toward instrumental harm, the latter captures impartial concern for the greater good. The

⁴ In Study 6 I used the term ‘cognitive capacities’ and did not include information about the individuals’ social emotions in the definition of that notion. By contrast, in all later studies I used the term ‘mental capacities’ and did include information about the individuals’ social emotions in the definition of that notion. For simplicity reasons, I use the term ‘mental capacities’ throughout the chapter. Furthermore, in Study 6 we used the terms “limited”, “high”, and “very high” to describe the different mental capacity levels. In all later studies I used the terms “basic”, “advanced”, and “extremely advanced” instead. For simplicity reasons, I always use the latter terminology throughout the chapter.

Speciesism Scale, OUS, and demographic questions were included in all studies but Study 3.4, which did not include the two scales.

Throughout all my studies, I tested whether participants would prioritise one individual over another through conducting a one-sample t-test against the mid-point. If responses were significantly above or below the mid-point, I concluded that participants tended to prioritise one individual over the other, and if responses did not significantly differ from the mid-point, I concluded that participants tended to be undecided and flip a coin to decide whom to prioritise.

Results

I first looked at the nine ‘inter-species’ dilemmas in which humans were pitted against monkeys. In total there were nine such dilemmas, of three different types: three ‘superior-humans’ dilemmas, where the humans had more advanced mental capacities than the monkeys (extremely advanced vs. advanced, extremely advanced vs. basic, advanced vs. basic), three ‘superior-monkeys’ dilemmas, where the monkeys had more advanced mental capacities (the same three pairs), and three ‘equal-capacities’ dilemmas, where the humans and monkeys had equal mental capacities (extremely advanced vs. extremely advanced, advanced vs. advanced, basic vs. basic). For each of the three types, the results from the three dilemmas converged. For example, there was no difference between the way a mentally advanced human was prioritised over a mentally basic monkey and the way a mentally extremely advanced human was prioritised over a mentally advanced monkey, $t(103) = -1.43, p = .15$. Hence, I created averaged scores for the three types and used them for all further analyses (*Figure 3.1*).

In line with my hypothesis, I found that in all cases where humans were pitted against monkeys, participants prioritised humans over monkeys, irrespective of the individuals’ mental capacity levels. Participants prioritised humans over monkeys in the superior-human dilemmas ($M = 2.24, SD = 1.57$), $t(103) = -11.39, p < .001, d = 1.12$, in the equal-capacities dilemmas ($M = 2.43, SD = 1.36$), $t(103) = -11.75, p < .001, d = 1.15$, as well as in the superior-monkey dilemmas ($M =$

2.91, $SD = 1.71$), $t(103) = -6.49$, $p < .001$, $d = .64$. A one-way ANOVA revealed significant differences across these three averaged score variables, $F(2, 309) = 5.14$, $p = .006$, $\eta_p^2 = .032$. Tukey's HSD post-hoc analyses showed that participants' tendency to prioritise humans over monkeys was significantly stronger when the humans' mental capacities were superior to the monkeys', than when the humans' mental capacities were inferior to the monkeys, $p = .006$. However, the differences in responses between the superior-human and equal-capacities dilemmas was not significant, $p = .65$, and neither were the differences in responses between the superior-monkey and equal-capacities dilemmas, $p = .07$.

Next, I looked at the six 'intra-species' dilemmas where individuals of the same species, but with different mental capacities, were pitted against each other (extremely advanced vs. advanced, extremely advanced vs. basic, advanced vs. basic). Again, I created two averaged scores composed of the three dilemmas involving humans, and the three dilemmas involving monkeys (Figure 3.2). Participants largely favoured flipping a coin when two humans with different mental capacities were pitted against each other ($M = 4.22$, $SD = 1.48$), $t(103) = 1.51$, $p = .14$, $d = .15$. However, participants consistently prioritised monkeys with superior mental capacities over monkeys with inferior mental capacities ($M = 4.96$, $SD = 1.56$), $t(103) = 6.23$, $p < .001$, $d = .61$. A paired t-test showed that participants were on average more likely to prioritise mentally superior monkeys over mentally inferior monkeys than they were to prioritise mentally superior humans over mentally inferior humans, $t(103) = -6.35$, $p < .001$, $d = .48$.

Since the correlations for each of the dilemmas were roughly the same, I created a single score for the tendency to prioritise humans over monkeys. I found that the more speciesist participants were, the more they tended to prioritise humans over monkeys ($r = -.34$, $p < .001$). Neither endorsement of instrumental harm ($r = .07$, $p = .51$) nor impartial beneficence ($r = .11$, $p = .51$) correlated with the tendency to prioritise humans over monkeys. However, the more participants endorsed instrumental harm, the more likely they were to prioritise humans ($r = .28$,

$p < .001$) and monkeys ($r = .22, p = .02$) with superior mental capacities over members of the same species with inferior mental capacities. Endorsement of impartial beneficence neither correlated with the tendency to prioritise between two humans with different mental capacities ($r = .13, p = .18$) nor with the tendency to prioritise between two monkeys with different mental capacities ($r = -.03, p = .80$). There were no correlations with demographic variables.





Figure 3.1. Participants prioritised humans over monkeys in all types of dilemmas: in those where the humans had superior mental capacities, in those where humans and monkeys had equal mental capacities, and in those where monkeys had superior mental capacities. (Study 3.1)

Figure 3.2. Participants did not prioritise mentally superior humans over mentally inferior humans, but they did prioritise mentally superior monkeys over mentally inferior monkeys. (Study 3.1)

Discussion

My results suggest that it is not the case that people value humans more than animals primarily because individual humans have advanced mental capacities, whereas individual animals have basic mental capacities. In *all* cases where humans were pitted against monkeys, participants

prioritised helping the human over the monkey. This was even true when the monkey had more advanced mental capacities than the human.

However, participants still gave *some* moral weight to more advanced mental capacities. Specifically, I found that they tended to prioritise more mentally advanced monkeys over less mentally advanced monkeys. For humans, by contrast, participants were more egalitarian and thus less likely to prioritise more mentally advanced humans over less mentally advanced humans.

To conclude, this study suggests that, while people grant higher moral status to individuals with more advanced mental capacities, species-membership trumps the effect of mental capacities. This conclusion is incompatible with the Individual Mental Capacity View, the notion that people prioritise humans over animals simply because of the more advanced mental capacities of individual humans. This means that there must be another reason for why people value humans more than animals.

One methodological limitation of this (and the following) study is that I had to rely on fantastic scenarios involving hypothetical beings. Participants have never encountered such individuals and had to make up their mind about them on the spot. While it was impossible to investigate the research questions of this study without using such fantastical scenarios, it is possible that responses to them are sensitive to small differences in presentation or phrasing. At the same time, such scenarios have the advantage of allowing us to uncover the factors that shape people's moral judgments rather than just measure conventional moral views.

Study 3.2. Species-Typical Mental Capacity Level

In Study 3.1, I found that differences in individual mental capacity cannot be the main reason why people value animals less than humans. Here, I test the Species-Typical Mental Capacity View; a view that has been defended by many prominent philosophers (see e.g. Cohen, 1986; Finnis, 1997; Scanlon, 1998). This hypothesis says that people's valuation of an individual is

based on the mental capacities of the *typical member of the individual's species* (rather than on the individual's *own* mental capacities). It has not been empirically tested before.

This hypothesis says that monkeys with advanced or extremely advanced mental capacities are valued less than the typical human because they belong to a species whose typical mental capacities are inferior to those of the typical human. Similarly, it says that humans whose mental capacities are atypically basic for humans are still valued as much as other humans, who have the species-typical level of mental capacity. Hence, this hypothesis says that such humans, too, are valued more than monkeys with advanced or extremely advanced mental capacities.

Study 3.2 was designed to test this hypothesis. My research question was: is people's valuation of individual members of a species based on the mental capacities of the typical member of that species? My hypothesis—which was pre-registered at <https://osf.io/whp35>—was that this is not the case. Instead, I hypothesised that participants would always prioritise humans over non-humans irrespective of the typical mental capacity level of the non-human species. In order to test this hypothesis, I asked participants to imagine hypothetical extraterrestrial species called Atlans, which were either described as typically having advanced mental capacities or, as typically having basic mental capacities. (I introduced these hypothetical extraterrestrial species because I thought that participants would find it more credible that the typical member of a yet unknown extraterrestrial species has advanced mental capacities, than that the typical member of an animal species has advanced mental capacities.) I hypothesised that in dilemmas when such beings are pitted against humans or each other, typical species-level would not make a difference: extraterrestrials with typically advanced and typically basic mental capacities would be treated similarly. I also included a third condition that contrasted monkeys to humans in order to investigate whether there would be differences in people's judgments between the conditions involving extraterrestrials and monkeys.

Method

Participants. I recruited 614 US American participants online via MTurk. They received \$0.72 in payment (in line with US minimum wage) for their participation. Fifty-five were excluded for failing a comprehension check, leaving a final sample of 559 people (278 female; $M_{age} = 39.2$, $SD_{age} = 11.96$). I expected a small to medium effect size. Power analysis showed that 576 participants were required to detect an effect size of $f = 0.13$ with an alpha of 0.05, power of 0.80, and 3 groups. I aimed to recruit 600 participants to account for any exclusions.

Materials and Procedure. I varied the type of non-human species that was pitted against a human (between-subjects) as well as the individual mental capacity levels of all beings (within-subjects) across six dilemmas: basic human vs. basic non-human, advanced human vs. advanced non-human, advanced human vs. basic non-human, basic human vs. advanced non-human, basic human vs. advanced human, basic non-human vs. advanced non-human. The non-human beneficiaries were either typically mentally advanced extraterrestrials, typically mentally basic extraterrestrials, or monkeys (which typically are mentally basic), depending on the condition.

The instructions first described the typical mental capacity level of the species, e.g. “*Like humans, typical Atlans can form advanced thoughts, have complex language, and can engage in cultural activities and politics (...)*”. Next, it was stated that there are a few special members of this species with more basic [advanced] mental capacities, e.g. “*However, while most typical Atlans have comparable mental capacities to humans, a few special Atlans you will encounter, who are severely mentally disabled, have basic mental capacities. Their mental capacities are much more basic than those of a typical Atlan or of a typical human and are comparable to those of a typical monkey. (...)*”.

Therefore, the beneficiaries were either typical or atypical representatives of the species they belonged to. This was crucial to test whether individuals with equal individual mental capacity levels were treated differently depending on their species’ typical mental capacity level. Depending on the condition a mentally basic extraterrestrial was either a typical representative of its species

(in the typically mentally basic extraterrestrial condition) or an atypical representative of its species (in the typically mentally advanced extraterrestrial condition).

Like in Study 3.1, for each dilemma participants were asked to indicate which of the two individuals they would prioritise helping if both required urgent help, but they could only help one. While participants in Study 3.1 were told that they could give a painkiller medicine to one individual in order to relieve them from their pain, I used a more general phrasing in this and the following studies, with the aim of reducing potential confounders: “*Sometimes we can only save the life, treat the illness, or relieve the pain of some but not of others. We then have to choose: whose life do we save, whose illness do we treat, and whose pain do we relieve? Suppose you were in a situation where you could only help one of these two individuals: (A) A human that has basic mental capacities; (B) An Atlan that has basic mental capacities*”. Participants had to indicate on a 7-point scale which individual they should prioritise from an ethical standpoint (1 = *Definitely A*, 4 = *Flip a coin to decide*, 7 = *Definitely B*).

Results

For each of the six dilemmas, I conducted a one-way ANOVA to test whether there were significant differences across the three conditions (*Table 3.1*). I found that for all dilemmas there were significant differences across the conditions except—unsurprisingly—for the dilemma where a human was pitted against another human in all three conditions (thus, this dilemma was identical across the three conditions). Tukey HSD post-hoc analyses showed that none of the responses between the two extraterrestrial conditions differed significantly from each other (see *Table 3.1* for descriptive and test statistics). This means that, in line with my main hypothesis, it made no difference whether the extraterrestrials belonged to a typically mentally advanced or typically mentally basic extraterrestrial species. This was true both regarding dilemmas where an extraterrestrial was pitted against a human and regarding dilemmas where two extraterrestrials were pitted against each other. However, the responses in the conditions where monkeys were pitted against humans differed significantly from the responses in the two conditions where

extraterrestrials were pitted against humans (*Table 3.1*). The tendency to prioritise humans over non-humans was stronger when the non-human was a monkey than when it was an extraterrestrial.

In line with the results from Study 3.1, participants tended to prioritise humans over non-humans and to prioritise mentally more advanced over less advanced individuals (*Table 3.1* and *Figure 3.3*). For each condition (the monkey condition, the typically mentally advanced extraterrestrial condition, and the typically mentally basic extraterrestrial condition), I created averaged scores for the four dilemmas that pitted humans against non-humans. Participants prioritised humans over monkeys, $t(190) = -20.90, p < .001, d = 1.51$. Participants also prioritised humans over typically mentally advanced extraterrestrials, $t(186) = -8.85, p < .001, d = 0.65$. An exception, though, was the case where the individual extraterrestrial had advanced and the human basic mental capacities, in which case participants tended to flip a coin, $t(186) = 0.18, p = .86, d = .65$. Finally, it is worth noting that the variance across participants was large, as can be seen in *Figure 3.3*. Therefore, even though participants on average prioritised humans over non-humans, in several of the dilemmas many participants flipped a coin.

I conducted three paired-sample *t*-tests to test whether participants were more likely to prioritise mentally advanced non-humans over mentally basic non-humans than they were to prioritise mentally advanced humans over mentally basic humans (as they were in Study 3.1). That was indeed the case in all three conditions: the typically mentally advanced extraterrestrial condition, $t(186) = 3.39, p < .001, d = .25$; the typically mentally basic extraterrestrial condition, $t(180) = 4.71, p < .001, d = .35$; the monkey condition, $t(190) = 6.67, p < .001, d = .48$.

Like in Study 3.1, the more speciesist participants were, the more likely they were to prioritise humans over monkeys (averaged over all four dilemmas that involved different species; $r = -.28, p < .001$). Similarly, the more speciesist participants were, the more likely they were to prioritise humans over extraterrestrials if they were members of a typically basic extraterrestrial species ($r = -.27, p < .001$). However, speciesism did *not* correlate with the tendency to prioritise

humans over extraterrestrials if they were members of a typically advanced extraterrestrial species ($r = .01, p = .91$). A similar pattern was found for impartial beneficence. In contrast to Study 3.1, the more participants endorsed impartial beneficence, the less they tended to prioritise humans over monkeys ($r = .17, p < .001$) and humans over typically mentally basic extraterrestrials ($r = .27, p < .001$). However, endorsement of impartial beneficence did not correlate with a reduced tendency to prioritise humans over typically mentally advanced extraterrestrials ($r = .06, p = .42$). Endorsement of instrumental harm did not correlate with the tendency to prioritise humans over non-humans in any condition. But like in Study 3.1, endorsement of instrumental harm correlated with the tendency to prioritise mentally more advanced over less advanced beings of the same species, both for humans ($r = .15, p < .001$) and non-humans ($r = .11, p = .009$).

Table 3.1. Prioritisation judgments M (SD) of Study 3.2. 1 stands for prioritising the first individual, 4 for flipping a coin to decide whom to prioritise, and 7 for prioritising the second individual. Rows represent the dilemmas in which two individuals are pitted against each other. Columns represent the type of species the respective non-human being is a member of, i.e. the three between-subjects conditions.

	Non-human species			All three conditions	F	p	Typically advanced Atlans vs. typically basic Atlans	Typically advanced Atlans vs. monkeys	Typically basic Atlans vs. monkeys
	Typically advanced Atlans	Typically basic Atlans	Monkeys				p	p	p
Basic human vs. basic non-human	3.02 (1.39)	2.77 (1.50)	1.86 (1.40)	33.99	< .001	.22	< .001	< .001	
Advanced human vs. advanced non-human	3.09 (1.48)	2.94 (1.61)	2.06 (1.46)	25.68	< .001	.62	< .001	< .001	
Advanced human vs. basic non-human	2.73 (1.83)	2.43 (1.67)	1.81 (1.48)	15.2	< .001	.18	< .001	.001	
Basic human vs. advanced non-human	4.03 (2.04)	3.88 (1.93)	2.80 (1.97)	21.67	< .001	.76	< .001	< .001	

Basic human vs. advanced human	4.65 (1.89)	4.71 (1.74)	4.65 (1.68)	0.06	.95	.95	.99	.96
Basic non-human vs. advanced non-human	4.96 (1.81)	5.17 (1.62)	5.45 (1.70)	3.86	.02	.49	.02	.25




Figure 3.3. Participants tended to prioritise humans over non-humans. It made no difference whether the individual extraterrestrials (Atlans) were members of a typically mentally advanced or a typically mentally basic extraterrestrial species. The tendency to prioritise humans over extraterrestrials, however, was weaker than the tendency to prioritise humans over monkeys. For more detailed results of the individual dilemmas see *Table 3.1.* (Study 3.2)

Discussion

I conducted this study to test the Species-Typical Mental Capacity hypothesis: that people value humans more than animals because the typical human has advanced mental capacities, whereas the typical member of an animal species does not. The results speak against this hypothesis. First, participants prioritised humans over typically mentally advanced extraterrestrials, contradicting what the Species-Typical Mental Capacity Hypothesis predicts. Second, people treated members of typically mentally advanced and typically mentally basic extraterrestrials species equally, both when they were pitted against humans and against each other. This again contradicts

the Species-Typical Mental Capacity Hypothesis. I also found—in line with Study 3.1—that the *individual's* mental capacity *does* influence the moral status people attribute to it, even though species-membership has a much stronger influence.

Surprisingly, I found different correlational patterns in the typically mentally basic and typically mentally advanced extraterrestrials conditions. Speciesism correlated with prioritising humans over extraterrestrials in the typically mentally basic extraterrestrials condition (as well as the monkeys condition), but not in the typically mentally advanced extraterrestrials condition. Similarly, the more participants endorsed impartial beneficence, the less they tended to prioritise humans over members of typically mentally basic extraterrestrials (as well as monkeys), but not in the typically mentally advanced extraterrestrials condition. This suggests that some people may be sensitive to the species-typical mental capacity level in certain ways. However, given that participants prioritised humans over extraterrestrials to the same extent across the two extraterrestrial conditions, I still conclude that people in general do not substantially take the typical mental capacity level of a species into account when attributing moral status to its members. There were no correlations with demographic variables.

Another surprising finding was that people appear to value extraterrestrials more than monkeys. They are more inclined to prioritise humans over monkeys than to prioritise humans over extraterrestrials (including typically mentally basic extraterrestrials). I will investigate this finding further in Study 3.4.

Study 3.3a. Species-Relativism and Pro-Human Bias

In Study 3.1 and 3.2, I tested whether people value humans more than animals because of their respective individual or species-typical mental capacity level. I found that these hypotheses cannot fully explain why people value humans more than animals. This is because I found that people *always* give priority to humans over non-humans regardless of their respective individual

mental capacities, and regardless of the typical mental capacity level of the non-human species. While individual mental capacity made some difference, it could not explain the massive difference in moral status people ascribe to humans and non-humans. In other words, people exhibit speciesism. In Study 3.3a, I wanted to further investigate the source of this speciesism. One hypothesis is that people think that humans should value other humans more, because we belong to the same species (Species-Relativism). Alternatively, it could be that people think that humans should be prioritised not because of shared species-membership, but merely because they are humans (Pro-Human Bias).

To tease these two hypotheses apart I modify a factor which I have not varied in the previous studies, namely the species-membership of the would-be rescuer. In my previous studies, participants were asked how they themselves think they should decide in such a situation. This means that the rescuer was always a human being. In this study, by contrast, I ask participants how they think that rescuers belonging to different species should decide. Thus, the species-membership of the rescuer was manipulated. This allows us to tease apart the Species-Relativism and Pro-Human Bias hypotheses. The Species-Relativism hypothesis predicts that participants will think that the rescuer should always prioritise members of their own species over other individuals, including humans. By contrast, the Pro-Human Bias hypothesis predicts that participants will think that all rescuers, irrespective of species-membership, should always prioritise humans over other individuals, including members of the rescuer's own species.

Methods

Participants. I recruited 184 US-American participants online via MTurk. They received \$0.72 in payment (in line with US minimum wage) for their participation. Fifteen participants were excluded for failing at least one attention or comprehension check or not finishing the study, leaving a final sample of 169 people (73 female; $M_{age} = 40.30$, $SD_{age} = 12.33$). I expected a medium effect size. Power analysis showed that 159 participants were required to detect an effect size of f

= 0.25 with an alpha of 0.05, power of 0.80, and 3 groups. I aimed to recruit 170 participants to account for any exclusions.

Materials and Procedure. The study had three between-subjects conditions, where the species of the rescuer was varied: a human rescuer condition, an extraterrestrial (again called ‘Atlas’) rescuer condition and a monkey rescuer condition. Participants in each condition were presented with three moral dilemmas in randomised order, analogous to the ones in Study 3.3a. The three dilemmas pitted a human beneficiary against a monkey beneficiary, a human beneficiary against an extraterrestrial (Atlas) beneficiary, and a monkey beneficiary against an extraterrestrial beneficiary. In contrast to the previous two studies, the mental capacity levels of the beneficiaries were not systematically varied. Instead, the mental capacity levels of the beneficiaries were consistently described as advanced for humans and extraterrestrials, but consistently as basic for monkeys (since that is the mental capacity level that currently existing monkeys have). However, in the monkey rescuer condition, it was said that the monkey rescuer had advanced mental capacities (in contrast to the monkey beneficiary). I made this choice because I wanted to make clear that the monkey rescuer could engage in complex ethical reasoning. Again, participants were asked to indicate which of the two individuals the rescuer should prioritise from an ethical standpoint on a 7-point scale (1 = *Definitely A*, 4 = *Flip a coin to decide*, 7 = *Definitely B*). All instructions were identical to those in Study 3.2, apart from the fact that I now asked the individuals what another individual should do (rather than what they themselves should do).

Results

In all three dilemmas there were significant differences across the conditions (*Figure 3.4*). First, there were significant differences across the conditions in the human vs. monkey-dilemma, $F(2, 166) = 51.68, p < .001, \eta^2 = .38$. Participants thought that both the human ($M = 1.53, SD = 0.71; t(56) = 26.31, p < .001; d = 3.49$) and the extraterrestrial rescuer ($M = 1.88, SD = 1.25; t(56) = 12.78, p < .001, d = 1.69$) should prioritise the human over the monkey beneficiary. However,

participants thought that the monkey rescuer should flip a coin to decide between saving a human or a monkey beneficiary ($M = 4.00$, $SD = 1.94$; $t(54) = 0$, $p = 1$, $d = 0$).

Second, there were significant differences across the conditions in the human vs. extraterrestrial-dilemma, $F(2, 166) = 11.65$, $p < .001$, $\eta^2 = .12$. Participants thought that a human ($M = 2.53$, $SD = 1.27$; $t(56) = 8.77$, $p < .001$, $d = 1.16$) as well as a monkey rescuer ($M = 3.62$, $SD = 1.24$; $t(54) = -2.28$, $p = .03$, $d = .31$) should prioritise the human over the extraterrestrial beneficiary. However, they thought that an extraterrestrial rescuer ($M = 3.77$, $SD = 1.89$; $t(56) = -0.91$, $p = .37$, $d = .12$) should flip a coin to decide between saving a human or an extraterrestrial beneficiary.

Third, there were also significant differences across the conditions in the monkey vs. extraterrestrial-dilemma, $F(2, 166) = 32.57$, $p < .001$, $\eta^2 = .28$. Participants thought that both a human ($M = 5.86$, $SD = 1.36$; $t(56) = 10.36$, $p < .001$, $d = 1.37$) as well as an extraterrestrial rescuer ($M = 6.00$, $SD = 1.25$; $t(56) = 12.05$, $p < .001$, $d = 1.60$) should prioritise the extraterrestrial over the monkey beneficiary. However, participants thought that a monkey rescuer should flip a coin to decide between saving a monkey or an extraterrestrial beneficiary ($M = 3.93$, $SD = 1.87$; $t(54) = -0.29$, $p = .77$, $d = .04$).

The more speciesist participants were, the more they thought that a human rescuer should prioritise a human beneficiary over a monkey beneficiary ($r = -.37$, $p = .004$) and over an extraterrestrial beneficiary ($r = -.37$, $p = .006$). As expected, there was no correlation between speciesism and judgments of how a human rescuer should prioritise between a monkey and an extraterrestrial beneficiary ($r = .08$, $p = .54$). Furthermore, the more speciesist participants were, the more they thought that an extraterrestrial rescuer should prioritise a human beneficiary ($r = -.36$, $p = .005$) and an extraterrestrial beneficiary ($r = .28$, $p = .04$) over a monkey beneficiary. However, there was no significant correlation between speciesism and judgments of how an extraterrestrial rescuer should prioritise between an extraterrestrial and a human beneficiary ($r = -$

.25, $p = .06$). Finally, speciesism did not correlate with responses in the dilemmas in which the monkey was the rescuer.

I did not find relevant correlations between the two subscales of OUS—endorsement of instrumental harm and impartial beneficence—and responses in the dilemmas in any of the conditions. There were no relevant correlations with demographic variables.

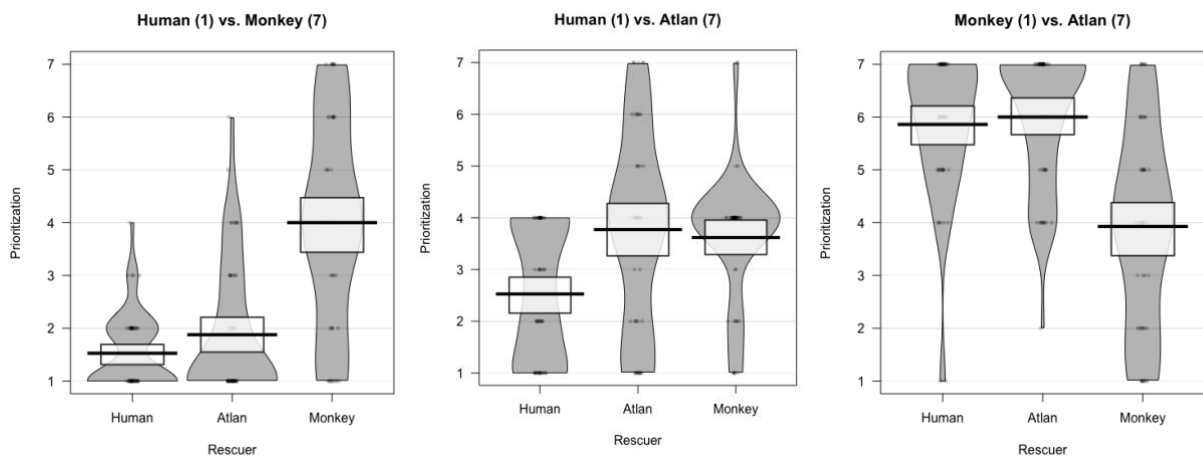


Figure 3.4. Participants tended to say that humans should prioritise themselves over others, that extraterrestrials (Atlans) should sometimes prioritise members of their own species over others, and that monkeys should flip a coin to decide who to prioritise (Study 3.3a).

Discussion

In Study 3.1, I found that while people do give greater moral priority to individuals with more advanced mental capacities over those with lower ones (Individual Mental Capacity View), this factor does not fully explain the massive moral priority that people give to humans over animals. I concluded that people prioritise humans over animals partly purely because of species-membership: they exhibit speciesism. In Study 3.3a I studied what underlies this phenomenon. I pitted two hypotheses against each other: Species-Relativism, the view that each species should prioritise their own, and Pro-Human Bias, the view that everyone, including non-humans, should prioritise humans. I found that neither of these hypotheses fully explained my findings. While the

Species-Relativism hypothesis found support in the fact that participants thought that extraterrestrials should prioritise members of their own species over monkeys, and in the fact that humans should prioritise humans over non-humans, it is contradicted by the fact that participants thought that monkeys should flip a coin to decide between helping a human or a monkey, and an extraterrestrial or a monkey. It is also contradicted by the finding that participants thought that an extraterrestrial should flip a coin to decide between helping a *human* and an extraterrestrial. This latter finding is of particular interest, since unlike in the human vs. monkey dilemma, in this dilemma both beneficiaries had the same mental capacities. The Pro-Human Bias hypothesis, in turn, was contradicted by the fact that participants did not think that non-humans should always prioritise humans over others.

How can these findings be reconciled? My interpretation is that, as hypothesised in the introduction, people prioritise humans over animals for a multiplicity of reasons. People give *some* moral weight to individual mental capacity, as well as to Species-Relativism and Pro-Human Bias. In the standard case where humans are pitted against animals with inferior mental capacities, all three factors are at work, leading to the familiar massive priority given to humans over other animals. But in other scenarios only one or two factors are at play. For instance, when people are asked how monkeys should prioritise between humans and monkeys, Species-Relativism is not at play, leaving only the Mental Capacity and Pro-Human Bias factors. When people are asked how intelligent extraterrestrials should prioritise between humans and a member of their own species, only the Pro-Human Bias factor is active. Accordingly, people prioritise humans less in those cases.

However, one problem with Study 3.3a was that several preferences were present simultaneously, which makes it difficult to tease them apart. In Study 3.3b, I address this issue by keeping the mental capacity levels of all beneficiaries constant, such that they only differ in their species-membership. This will allow us to more directly contrast the Species-Relativism and Pro-Human Bias views and eliminate confounding preferences.

Study 3.3b. Species-Relativism and Pro-Human Bias

In Study 3.3a, I showed that Species-Relativism and Pro-Human Bias likely are parts of the explanation of the moral anthropocentrism. In this study—which was pre-registered at <https://osf.io/3nf9y>—I aimed to replicate this finding in a more controlled setting. I wanted to study the effects of Species-Relativism and Pro-Human Bias without potential confounders such as a preference for individuals with more advanced mental capacities, or a possible preference for extraterrestrial over terrestrial non-humans. Therefore, I included only species with equally advanced mental capacities: humans, Atlans, and another hypothetical mentally advanced extraterrestrial species, called Bredlans. Participants were either asked what they think a human rescuer should do, or what an Atlan rescuer should do.

Several of my hypotheses pertained to the Species-Relativism view—that people believe that individuals should prioritise members of their own species. First, I hypothesised that participants would be more likely to say that the Atlan rescuer should prioritise the Atlan over the Bredlan than that the human rescuer should do that. Second, I hypothesised that participants would be more likely to say that the human rescuer should prioritise the human over the Bredlan than that the Atlan rescuer should do that. Third, I hypothesised that participants would be more likely to say that the human rescuer should prioritise the human over the Atlan than that the Atlan rescuer should do that. These hypotheses left open the possibility of Pro-Human Bias playing a role as well. For example, I hypothesised that Pro-Human Bias would lead people to find it more important for a human rescuer to prioritise a human over an Atlan than for an Atlan rescuer to prioritise an Atlan over a human. Finally, I hypothesised that religious people would be more likely to hold a Pro-Human Bias because several religions (e.g. Christianity) hold the view that humans are a special species that is in an absolute sense more valuable than animals. To test this, I also measured religiosity.

Methods

Participants. I recruited 225 US American participants online via MTurk. They received \$0.6 in payment (in line with US minimum wage) for their participation. Eleven were excluded for failing at least one attention or comprehension check or for not completing the study, leaving a final sample of 214 people (100 female; $M_{age} = 40.65$, $SD_{age} = 11.75$). I relied on a small to medium effect because I expected there to be only a small difference between the conditions. Power analysis showed that I would need a total sample of 200 to detect a small to medium effect of $d = .4$, relying on an alpha of 0.05, power of 0.8 (two-tailed). I aimed to recruit 220 participants to account for exclusions.

Materials and Procedure. The study design was based on Study 3.3a, though in study 3.3b I only had two between-subjects conditions: one where the rescuer was an Atlan, and one where the rescuer was a human. The beneficiaries were either humans, Atlans, or Bredlans. Humans, Atlans and Bredlans were described as being different species with similarly advanced mental capacity levels. Apart from these differences, instructions were identical to those of Study 3.3a. Again, participants were presented with three moral dilemmas in randomised order. The three dilemmas pitted an Atlan against a Bredlan beneficiary, a human against an Atlan beneficiary, and a human against a Bredlan beneficiary. Again, participants were asked to indicate which of the two individuals the rescuer should prioritise from an ethical standpoint on a 7-point scale (1 = *Definitely A*, 4 = *Flip a coin to decide*, 7 = *Definitely B*).

After the main task, participants were asked to indicate their agreement with three explicit moral principles on 7-point Likert scales (1 = *Strongly disagree*, 7 = *Strongly agree*): the egalitarian aspect of the Mental Capacity view (“*Beings with equal mental capacities should matter equally*”), Species-Relativism (“*Beings should prioritise members of their own species over members of other species*”), and Pro-Human Bias (“*Humans are always morally more important than other types of beings*”). Finally, they responded to the Centrality of Religiosity Scale (CRS; Huber & Huber, 2012).

Results

All three dilemmas yielded significant differences (*Figure 3.5*), and there were no effects of the order of the dilemmas on the participants' responses. First, as hypothesised, participants considered it more important for the Atlan rescuer to prioritise an Atlan over a Bredlan beneficiary ($M = 3.24, SD = 1.39$) than for a human rescuer to do so ($M = 3.98, SD = 0.37$), $t(125.1) = 5.43, p < .001, d = .72$. Participants thought that the Atlan should clearly prioritise the Atlan, $t(109) = 5.77, p < .001, d = .55$, whereas the human rescuer should flip a coin to decide between prioritising an Atlan or Bredlan, $t(103) = 0.53, p = .60, d = .05$.

Second, as hypothesised, participants considered it more important for the human rescuer to prioritise a human over a Bredlan ($M = 3.02, SD = 1.22$) than for an Atlan rescuer to do so ($M = 3.74, SD = 1.22$), $t(211.2) = 4.30, p < .001, d = .59$. Participants thought that both the human ($t(103) = 8.18, p < .001, d = .80$) and the Atlan rescuer ($t(109) = 2.27, p = .03, d = .22$) should prioritise a human over a Bredlan beneficiary. However, note that the effect size for the Atlan rescuer was small and that almost all participants chose to flip a coin in this condition.

Third, as hypothesised, participants thought that a human rescuer should prioritise a human over an Atlan ($M = 2.85, SD = 1.25$), $t(103) = 9.40, p < .001, d = .92$. Further, I found that an Atlan rather should prioritise an Atlan over a human ($M = 4.53, SD = 1.58$), $t(109) = 3.50, p < .001, d = .33$. Crucially, participants were more likely to say that a human should prioritise a human over an Atlan, than to say that an Atlan should prioritise an Atlan over a human, $t(206) = 3.22, p = .002, d = .44$.




Figure 3.5. Participants tended to say that both humans and intelligent extraterrestrials should prioritise members of their own species over other individuals. In addition, there was a slight overall tendency towards prioritising humans. (Study 3.3b)

I found several correlations between explicitly endorsed moral principles and participants' responses in the dilemmas (Table 3.2). For example, I found correlations between endorsing the Pro-Human Bias principle and consistently prioritising humans and between endorsing the Species-Relativism principle and saying that humans as well as Atlans should prioritise members of their own species. I also found that endorsement of the Mental Capacity principle correlated negatively with endorsement of Species-Relativism ($r = -.37, p < .001$) as well as with endorsement of Pro-Human Bias ($r = -.38, p < .001$), whereas endorsement of Species-Relativism and endorsement of Pro-Human Bias correlated positively with each other ($r = .66, p < .001$).

Table 3.2. Correlations between the three explicitly endorsed moral principles and prioritisation decisions.

	Mental Capacity principle <i>M</i> = 5.53, <i>SD</i> = 1.36		Species-Relativism principle <i>M</i> = 4.17, <i>SD</i> = 1.89		Pro-Human Bias principle <i>M</i> = 3.26, <i>SD</i> = 1.86	
	Human rescuer	Atlan rescuer	Human rescuer	Atlan rescuer	Human rescuer	Atlan rescuer
Atlan vs. Bredlan beneficiary	.10	.10	-.05	-.57***	-.14	-.16
Human vs. Atlan beneficiary	.48***	.11	-.81***	.37***	-.63***	-.06
Human vs. Bredlan beneficiary	.47***	.02	-.73***	-.17	-.52***	-.31***

Positive correlations mean that the stronger the endorsement of the principle, the more the latter individual is prioritised over the former and vice versa. Mental Capacity principle (egalitarian aspect): “*Beings with equal mental capacities should matter equally*”; Species-Relativism principle: “*Beings should prioritise members of their own species over members of other species*”; Pro-Human Bias principle: “*Humans are always morally more important than other types of beings*”

The more speciesist participants were, the more they tended to say that a human should prioritise a human over an Atlan ($r = -.36, p < .001$) and a Bredlan ($r = -.23, p = .02$). As expected,

Speciesism did not correlate with participants' responses when the human rescuer had to decide between an Atlan and Bredlan ($r = -.12, p = .23$). Neither were there any significant correlations between speciesism and dilemma responses when the rescuer was an Atlan. Speciesism correlated positively with explicit endorsement of both the Species-Relativism ($r = .36, p < .001$) and Pro-Human Bias ($r = .49, p < .001$) principles, but negatively with the egalitarian aspect of the Mental Capacity view ($r = -.28, p < .001$).

The more participants endorsed instrumental harm (OUS), the more they thought that a human rescuer should prioritise a human over an Atlan ($r = -.21, p = .03$) and a Bredlan ($r = -.22, p = .02$). Similarly, the more they endorsed instrumental harm, the more they thought that an Atlan rescuer should prioritise an Atlan over a Bredlan ($r = -.28, p = .003$). However, stronger endorsement of instrumental harm did not correlate with thinking that an Atlan rescuer should prioritise an Atlan over a human ($r = -.08, p = .41$). Similarly, endorsement of instrumental harm did not correlate with responses in the dilemma where a human rescuer had to decide between saving an Atlan and Bredlan ($r = -.14, p = .15$) or where an Atlan rescuer had to decide between saving a human and a Bredlan ($r = -.08, p = .42$). The more participants endorsed instrumental harm, the more they endorsed the Species-Relativism principle ($r = .36, p < .001$) and the Pro-Human Bias principle ($r = .32, p < .001$), and the less they endorsed the egalitarian Mental Capacity principle ($r = -.19, p < .001$). The more strongly participants endorsed impartial beneficence (OUS), the less they thought that the human rescuer should prioritise the human over the Atlan ($r = .28, p = .004$) and the Bredlan ($r = .32, p = .001$). However, there were no further significant correlations between endorsement of impartial beneficence and dilemma responses in either of the two conditions. Endorsement of impartial beneficence correlated negatively with endorsement of the Species-Relativism principle ($r = -.22, p < .001$), positively with endorsement of the Mental Capacity principle ($r = .30, p < .001$), and not significantly with endorsement of the Pro-Human Bias principle ($r = -.08, p = .25$).

The more politically conservative participants were, the more they thought that a human rescuer should prioritise a human over an Atlan ($r = -.29, p = .003$) and a Bredlan ($r = -.24, p = .01$), and similarly, the more they thought that an Atlan rescuer should prioritise an Atlan over a Bredlan ($r = -.20, p = .04$). In line with this, the more politically conservative participants were, the more they endorsed the Species-Relativism principle ($r = .37, p < .001$) and the Pro-Human Bias principle ($r = .44, p < .001$), but the less they endorsed the Mental Capacity principle ($r = -.25, p < .001$). There were no other correlations with political orientation.

Against my hypothesis, there were no significant associations between religiosity—as measured by both CRS ($\alpha = .93$) as well as by stated religious affiliations such as Catholicism—and dilemma responses. There were no further noteworthy correlations between demographic variables.

Discussion

The results of this study further confirm the existence of Species-Relativism. People believe that individuals should prioritise members of their own species over members of other species. Participants believed that humans should prioritise other humans over non-humans and that intelligent extraterrestrials (Atlans) should prioritise members of their own species over other members of other species of intelligent extraterrestrials (Bredlans), and even over humans.

Furthermore, this study also supports the existence of Pro-Human Bias. This is shown in the (weak) preference participants had in favour of humans. For example, they considered it more important for a human to prioritise another human over Atlans than for an Atlan to prioritise another Atlain over humans. Also, participants tended to believe that an Atlan should prioritise a human over a Bredlan. Note, though, that these are statements about the average tendency and that it may be that only a small minority of participants show a Pro-Human Bias, whereas most do not. For example, 72% of participants believed an Atlan should flip a coin to decide between

prioritising a human and a Bredlan. Contrary to my hypothesis, religiosity was not associated with Pro-Human Bias (nor with Species-Relativism).

The hypothesis that both Species-Relativism and Pro-Human Bias both contribute to moral anthropocentrism is supported not only by the dilemma responses, but also by the finding that some people endorse these views as explicit principles. Endorsement of the Species-Relativism principle was higher than that of Pro-Human Bias. Endorsement of the explicit principles (including the Mental Capacity view) correlated with the dilemma decisions in predictable ways. For example, the more participants endorsed the Species-Relativism principle, the stronger was their tendency to say that humans as well as Atlans should prioritise themselves over other species. And the more participants endorsed the Pro-Human Bias principle, the stronger was their tendency to say that both humans and Atlans should prioritise humans over non-humans. It is possible, though, that endorsement of the Pro-Human Bias principle was inflated due its ambiguity. It was perhaps not made clear enough that even non-human beings would have to agree that humans are the most valuable species. Instead, some participants may have interpreted the statement as the species-relative statement that they themselves as humans consider humans the most valuable species.

Study 3.4. Harming and Anti-Animal Bias

In all studies of this chapter so far, I presented participants with difficult prioritisation decisions, where they chose between helping a human or helping a non-human. These are decisions where, in other contexts, partiality is often seen as justified. For instance, most people would find it acceptable to prioritise helping one's own family members over others (Petrinovich et al., 1993). In cases where one party would be actively harmed for the sake of the other party, partiality is likely to be more controversial. However, with respect to animals, people do consider it relatively permissible to actively harm and even systematically exploit animals to benefit humans

as shown in the real-world context of consumption and medical experimentation as well as by recent psychological research (Chapter 2). I therefore wanted to investigate whether the factors that were found to drive moral anthropocentrism in the helping context also explain this permissive attitude to harm to animals. In Chapter 4, I will explore people's moral judgments about active harm towards animals in greater detail.

This attitude cannot be entirely explained by the lower mental capacities of animals (the Mental Capacity View) given that we do not accept such harm to humans with significant cognitive impairments. It is also unlikely to be fully explained by a combination of lower mental capacities and membership of a different species (Species-Relativism), given that in Study 3.2 I found that people gave greater priority to extraterrestrials with basic mental capacities than to animals with the same capacities. I therefore hypothesised that a further factor is at work here—a specific bias against beings that are categorised as animals, over and above their specific or typical capacities—which I call “Anti-Animal Bias”. However, an alternative explanation of that finding is that people perceived the hypothetical extraterrestrials as more human-like, or as having more advanced mental capacities, than the monkeys. In this study I therefore aimed to study whether people manifest such an Anti-Animal bias, and whether it may explain permissive attitudes to harming animals.

I hypothesised that participants—like in Study 3.2—would be more willing to not help a monkey than an Atlan with similar mental capacities, and that they would be more willing to actively harm a monkey than an Atlan in order to benefit a human. I further hypothesised that the difference in willingness to harm an animal compared to an Atlan would be greater than the difference in willingness to not help them. I assumed this based on the observation that harm towards and exploitation of animals in the real world is considered permissible to particularly strong extent.

Methods

Participants. I recruited 271 US American participants online via MTurk. They received \$0.6 in payment (in line with US minimum wage) for their participation. Twenty-two were excluded for failing at least one attention or comprehension check or for not completing the study, leaving a final sample of 249 people (109 female; $M_{age} = 42.50$, $SD_{age} = 12.83$). Based on the difference I found between typically mentally basic Atlans and monkeys in Study 3.2, I expected a small to medium effect size. Power analysis showed that I would need a total sample of 245 to detect an effect of $f = .18$, relying on an alpha of 0.05, power of 0.8, df of 1, and four groups. I therefore aimed to recruit 270 participants to account for exclusions.

Materials and Procedure. The study had a between-subjects 2 species (Atlan vs. monkey) x 2 type (help vs. harm) design. The vignettes were based on those of Study 3.2. Participants were informed that while most typical Atlans (or monkeys) had basic mental capacities, a few special Atlans (or monkeys) had mentally advanced mental capacities. Next, they were presented with two moral dilemmas in randomised order. In one of the dilemmas, a mentally basic Atlan (or monkey) was pitted against a mentally basic human, and in the other dilemma a mentally advanced Atlan (or monkey) was pitted against a mentally advanced human. Depending on the condition, the dilemmas either involved a helping situation identical to the ones in the previous four studies, or it involved a situation in which an Atlan (or monkey) had to be harmed in order to help a human. The harm dilemma read as follows: “*Suppose you were in a situation in which the only way to help a human that has basic mental capacities is by actively harming an Atlan [a monkey] that has basic mental capacities.*” Participants were then asked to indicate what they should do from ethical standpoint on a 7-point scale (1 = *Definitely NOT harm the Atlan [monkey]*, 4 = *Flip a coin to decide*, 7 = *Definitely harm the Atlan [monkey]*).

After the main task, participants responded to three follow-up questions. First, they were asked to indicate on a scale from 1 to 100 how they would rate the mental capacities of a typical

Atlan (or monkey) if a typical human was 100. Second, they were asked to indicate on a scale from 1 to 100 how similar a typical Atlan (or monkey) is to a human. Third, they were asked to indicate on a scale from 1 (*Not at all*) to 7 (*Completely*) to what extent they think a typical Atlan (or monkey) is an animal. These three questions were included to test whether any differences in the moral dilemma responses could be explained by either different perceptions of mental capacity levels, human-likeness or animal-likeness (our hypothesis).

Results

I created an average score for the responses of the two dilemmas (basic and advanced mental capacities) and ran a 2x2 between-subjects ANOVA, which yielded two main effects and no interaction (Figure 3.6). Participants were more likely to not help and to harm monkeys than Atlans, $F(1, 245) = 15.34, p < .001, \eta_p^2 = .06$. And they were more likely to deprioritise than to harm both types of beings, $F(1, 245) = 41.21, p < .001, \eta_p^2 = .15$. There was no interaction between species and type, $F(1, 245) = 0.37, p = .55, \eta_p^2 = .001$. The same type of results was found when ANOVAs were conducted on the responses of the two dilemmas individually (Table 3.3 for descriptive statistics).

Table 3.3. Means and standard deviations of the responses in the moral dilemmas involving mentally basic and mentally advanced individuals (Study 3.4). Lower scores indicate greater willingness to help humans instead of non-humans and greater willingness to actively harm non-humans to save humans (reversed).

	Atlan		Monkey	
	Harm	Help	Harm	Help
Basic	4.23 (2.05)	2.65 (1.56)	3.17 (1.97)	2.14 (1.50)
Advanced	4.46 (2.10)	2.97 (1.56)	3.56 (1.99)	2.05 (1.43)
Averaged	4.34 (2.00)	2.81 (1.45)	3.36 (1.88)	2.09 (1.42)

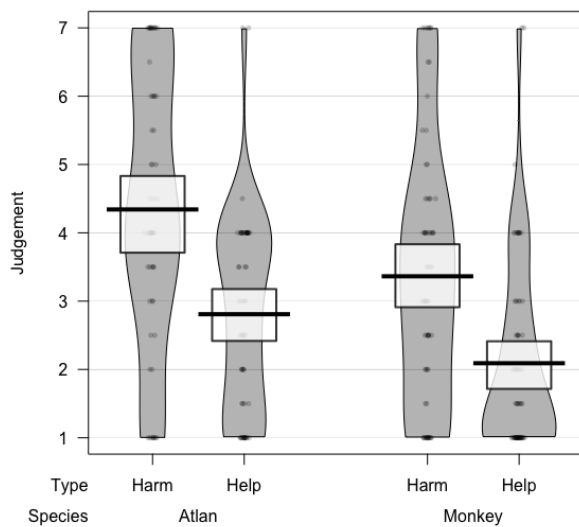


Figure 3.6. Participants were more likely to both not help as well as actively harm monkeys than Atlans (averaged scores). For both species, they were more willing to not help them than to actively harm them. In the harm conditions 1 stands for *Definitely harm the Atlan* and 7 for *Definitely NOT harm the Atlan* (reversed). In the help conditions 1 stands for *Definitely prioritise the*

human and 7 for *Definitely prioritise the Atlan*. 4 stands for *Flip a coin to decide* in all conditions (Study 3.4).

There were no significant differences in the level of mental capacities they believed typical Atlans ($M = 43.25$, $SD = 24.50$) had compared to typical monkeys ($M = 39.56$, $SD = 22.39$), $t(236.5) = 1.23$, $p = 0.22$. Similarly, there were no significant differences in the extent to which participants perceived Atlans ($M = 44.71$, $SD = 26.35$) and monkeys ($M = 45.60$, $SD = 24.88$) to be human-like, $t(239.4) = -0.27$, $p = 0.79$. However, participants perceived a typical monkey ($M = 6.50$, $SD = 1.04$) as an animal to a greater extent than they did a typical Atlan ($M = 4.28$, $SD = 1.85$), $t(177.7) = 11.48$, $p < .001$. There were no correlations with demographic variables.

Discussion

In line with my main hypothesis, this study demonstrated that people appear to have a specific bias against animals, which I call Anti-Animal Bias. Participants were more likely to deprioritise as well as harm monkeys than the hypothetical extraterrestrial Atlans despite the fact that the mental capacities of both were described in identical terms. Moreover, Atlans and monkeys were perceived as having similar levels of mental capacities and as being similarly human-like. Instead, the tendency to prioritise Atlans over monkeys—as far as I can conclude from this study—is purely a result of the fact that people perceived monkeys to be animals and Atlans less so.

Participants exhibited similar levels of priority to Atlans over monkeys to in help and harm scenarios. This suggests that there may be only a difference of degree between the helping judgments investigated in the previous studies and parallel judgments about harming. It is worth mentioning that there was an asymmetry between the harm and help scenarios I used. While in the help scenarios both options involved an action (to help one but not another one), in the harm

scenarios one option involved an action and the other an omission. If in the harm scenarios instead both options involved either harming a human or a non-human, people would presumably prioritise the human much stronger than in my study.

General Discussion

Why do people prioritise humans morally over animals? Both critics and supporters of moral anthropocentrism often see it as driven by a single core factor. By contrast, my studies provide strong evidence that there are multiple overlapping factors that underpin moral anthropocentrism in the general population.

The Mental Capacity View

If the only reason why people value humans more was their more advanced mental capacities, we should expect them to value non-humans with similarly or more advanced mental capacities as much as humans (Individual Mental Capacity View). However, that was not the case. In Study 3.1 I found that participants prioritised humans over animals *even* if animals had equal (or even *superior*) mental capacities. However, while manipulating individual mental capacity could never flip the priority from humans to animals, it did contribute to the *strength* to which participants prioritised humans over animals: the more advanced the mental capacities of the animals were, the weaker the extent to which participants prioritised humans over animals. Participants also prioritised mentally more advanced non-humans over mentally less advanced non-humans of the same species. These findings extend previous research that found associations between an animal's perceived mental capacity level and their perceived moral status (e.g. Bastian et al., 2012; Chapter 2). And crucially, these findings demonstrate that individual mental capacity *does* play a role in explaining moral anthropocentrism even if it cannot explain it in full.

My findings do rule out one important candidate for explaining moral anthropocentrism, the Species-Typical Mental Capacity View: the view that the moral status of an individual is based on the mental capacity level of the typical member of the species. Moral philosophers often invoke this view to explain why so-called ‘marginal cases’, such as severely cognitively impaired humans, *should* be given full moral status and treated radically differently than animals, even when they do not possess greater mental capacities (see e.g. Cohen, 1986; Finnis, 1997; Scanlon, 1998). However, in Study 3.2 I found that the typical mental capacity level of a species had no effect on participants’ judgments in the moral dilemmas. Of course, these results do not directly show that the Species-Typical Mental Capacity View is mistaken as an *ethical* claim. But they do reveal a radical disconnect between a popular defence of moral anthropocentrism by some philosophers and the factors that actually underpin the moral thinking of lay people.

The Sub-Factors of Speciesism

In line with my findings in Chapter 2, I found that people are speciesists—that they prioritise humans over animals simply because they are humans. I deepened and extended that research through identifying three forms of speciesism: Species-Relativism, Pro-Human Bias, and Anti-Animal Bias; the latter two of which are instances of Species-Absolutism.

First, studies 3.3a and 3.3b showed that people follow a Species-Relativist principle, according to which individuals of all species should prioritise members of their own species over others. For example, participants judged that mentally advanced extraterrestrial beings should prioritise members of their own species over mentally advanced members of other species, including humans. Participants were less likely, however, to accept that monkeys should prioritise themselves over more advanced species. This could be because they believe that mentally advanced beings are more valuable than mentally basic ones, which in this case counteracts Species-

Relativism. Alternatively, it is possible that people are less inclined to apply Species-Relativism to mentally basic species or simply to animals due to Anti-Animal Bias.

Species-Relativism appears to function in a way akin to partiality towards one's family. People generally prioritise members of their family over others (Petrinovich et al., 1993) while also holding that other people should similarly prioritise members of their families (Pugh et al., 2013). Patriotism is another example. Species-Relativism is analogous to a US American saying "America first" and at the same time considering it legitimate for a Frenchman to say "France first". Partiality to one's family members or compatriots are widely seen as morally permissible and, in some contexts, even obligatory (Keller, 2013; Miller, 2005). It has even been found that a significant number of people would prioritise their own pet over a foreign tourist, suggesting that partiality towards family (in this case, a pet) can even overturn anthropocentrism (Topolski et al., 2013). Importantly, prioritization based in partiality need not reflect a difference in moral status or value. Parents may prioritise their children over others without believing that their children have a higher moral status.

In addition to Species-Relativism I found evidence of tendencies that I call Species-Absolutism, since they apply irrespective of the relationship between the helper and beneficiary. The first Species-Absolutist tendency I have identified is a Pro-Human Bias. Most notably, participants were more likely to say that a human should prioritise a human over an Atlan than they were to say that an Atlan should prioritise an Atlan over a human. Pro-Human Bias differs from the "America first" patriotism example above but would be more like American Exceptionalism—seeing the US as superior to all other nations in an absolute sense. It is possible that Pro-Human Bias is primarily implicit and that most people would not endorse under reflection (e.g. when the bias is made explicit). Whether people—and which people—really believe humans are more valuable qua humans needs to be investigated further. A further question is how people would justify such a bias, given that it may operate even when the non-humans had identical or

superior cognitive capacities. My hypothesis that religious people might be more likely to show Pro-Human Bias—given that some religious traditions portray humanity as having special cosmic importance—was not confirmed.

The second Species-Absolutist tendency I have identified is an Anti-Animal Bias. In Study 3.4 I found that participants were more likely to deprioritise and harm monkeys than equally mentally capable Atlans, and this difference seems to be entirely driven by the monkeys being perceived as animals to a greater extent than the Atlans.

The finding that speciesism has multiple sub-factors appears at first glance inconsistent with the finding that speciesism is a one-dimensional construct, as revealed by the factor analysis of Chapter 2. To understand this apparent inconsistency, it is important to consider two differences between the approaches of Chapter 2 and Chapter 3. First, the purpose of the development of the Speciesism Scale in Chapter 2 was to measure the relatively course-grained general tendency to value humans more than animals in different realistic contexts. Accordingly, the items of the initial item pool did not differentiate between the relatively subtle differences between Species-Relativism, Human Supremacism and Anti-Animal Bias. The purpose of Chapter 3, in contrast, was to explicitly identify the potential sub-factors of this overall tendency, even if these sub-factors usually all act together. Second, while Chapter 2 focused on individual differences between participants, Chapter 3 focused on general tendencies. It is likely that individual differences between the identified sub-factors of speciesism are highly correlated. For example, presumably almost everyone who shows Human Supremacism also shows Anti-Animal Bias. While, in principle the different factors could diverge from each other in certain special cases, such as in the ones tested in Studies 3.3a and 3.3b, in reality such special cases are extremely rare. Therefore, Chapter 2 demonstrated that from a course-grained perspective, when observing the tendency in the real world, speciesism is a one-dimensional construct. The deeper investigation of

Chapter 3, however, revealed that fundamentally there are multiple factors to speciesism that usually act together.

Relations Among the Factors of Moral Anthropocentrism

Speciesism and the Individual Mental Capacity view both contribute to moral anthropocentrism, as we have seen. Let us now look more precisely at the causal relations between these factors (*Figure 3.7*). When someone attributes a lower moral value to one being (an animal) than to another (a human) because of speciesism, there are two judgments involved: the perception that the two beings belong to different species, and the moral judgment that one of the species is morally more valuable than the other. It is the combination of these two judgments that contribute to moral anthropocentrism (1). Similarly, when someone attributes a lower moral value to one being than to another because of the Individual Mental Capacity View, there are two judgments involved: the perception that one of the beings is mentally more advanced than the other, and the moral judgment that it therefore is morally more valuable (2).

These two links (1 and 2) are direct relations between the main factors I have demonstrated and moral anthropocentrism. However, there are also indirect causal links between these factors. First, there is an indirect link between speciesism and perceptions of relative mental capacities of humans and animals: speciesism causes people to attribute lower mental capacities to animals (3). This tendency has been referred to as de-mentalization and has been demonstrated in previous research that showed that people attribute lower mental capacities to food animals compared with non-food animals (e.g. Bastian et al., 2012). Further, in Chapter 2 I found that speciesism correlates with attributions of lower mental capacities to animals (which of course is only a correlational finding and leaves open whether the link is causal).

Second, there is an indirect link between speciesism and the Individual Mental Capacity view. People apply the Individual Mental Capacity View selectively to animals and not to humans (4). In studies 3.1 and 3.2 I found that participants valued animals who were mentally more

advanced more than less advanced animals. At the same time, however, participants did not as clearly value mentally more advanced humans more than mentally less advanced humans. This tendency of selectively applying the Individual Mental Capacity View only to animals could itself be considered speciesist.

Therefore, according to this framework, speciesism reduces the perceived moral status of an animal in multiple ways: both directly and indirectly. This suggests that speciesism plays a more central role in explaining moral anthropocentrism than one might think at first glance. Future research needs to verify and refine this framework in more detail. In particular, the causal involvement of the sub-factors of speciesism remain unclear.

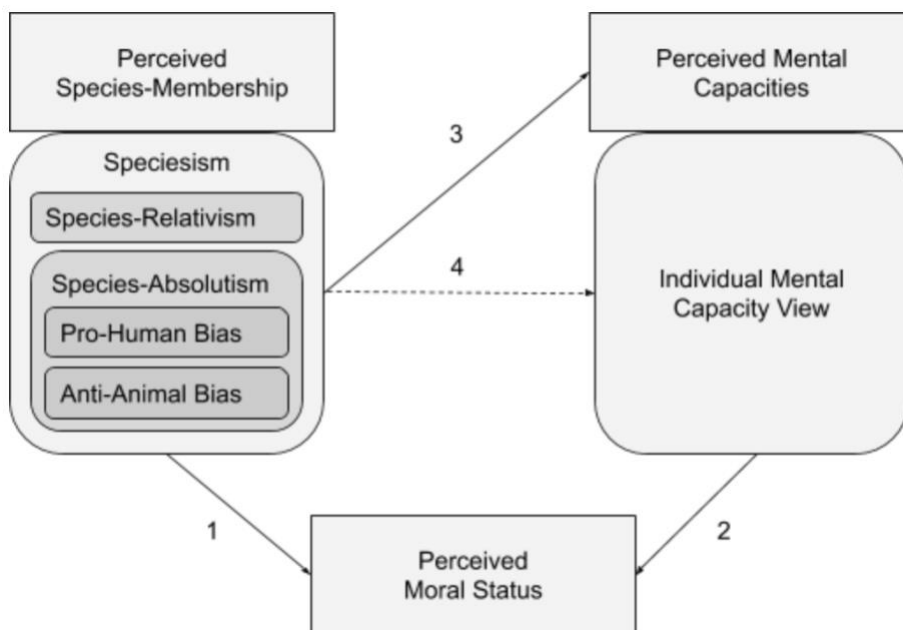


Figure 3.7. A Framework for Moral Anthropocentrism. One reason people perceive animals to have low moral status because of mere species-membership, i.e. speciesism (1). Another reason is the fact that people attribute lower moral status to beings that have low perceived mental capacities, i.e. Individual Mental Capacity View (2). Furthermore, there are two indirect paths from speciesism to moral anthropocentrism. First, speciesism leads people to have lower perceived mental capacities of animals (3). Second, due to speciesism, people selectively apply the Individual

Mental Capacity View only to animals but not humans (4), a phenomenon which is itself speciesist (i.e., speciesism appears to moderate the Individual Mental Capacity View).

The Sources of Moral Anthropocentrism

With the current evidence, many questions about the sources and causal role of the sub-factors of moral anthropocentrism and speciesism remain open. Here, I will make some speculations that future research could test.

The prioritization of humans over others could in significant part be an extension of familiar ways in which we prioritise ourselves and those linked to us, regardless of other considerations (e.g. to do with mental capacity). (In Chapter 5 I discuss the hypothesis that speciesism is a form of ingroup-favouritism in more detail.) Spelled out explicitly, as a ‘universalizable’ moral principle, this would generate the idea of Species-Relativism. But, perhaps less explicitly and consciously, it is also reflected in Pro-Human Bias—a degree of priority given to humans simply because they are human, independently of the actual relationship between helper and potential beneficiary.

In addition, people also have a desire to signal that they are fair moral agents, a disposition that is evolutionarily adaptive because it increases the chances of others engaging in mutual cooperation with us (DeScioli & Kurzban, 2013; Tooby & Cosmides, 2010). Processes of prioritization of the kind investigated here provide a context in which such fairness can be signalled. Accepting that others are permitted to prioritise their own group (in this case, a non-human species) can be a way to signal such fairness while at the same time having a justification to prioritise our own when we can. At the same time, since our expected audience is invariably human, treating humans as more important in a more absolute sense can still send an attractive signal—and possibly even more so in cultural contexts where contemporary ideas of moral equality are less dominant.

Crucially, though, people only have to signal their positive moral character traits to individuals who are themselves capable of engaging in moral reasoning and in potential mutual cooperation, i.e. to mentally advanced individuals. There is no gain in justifying to mentally basic individuals why they are being deprioritised because they would not understand the signal. This, in addition to similarity to self, can explain why people value mentally basic individuals less (in an absolute sense) than mentally advanced individuals (viz. the Individual Mental Capacity View). However, people are still capable of empathetic responses to the welfare of non-humans with lower mental capacities, and thereby still perceive them as having *some* moral status. Moreover, attitudes to individuals with lower mental capacities (typically, non-human animals) still send a valuable signal to other individuals with higher mental capacities—someone engaged in gratuitous cruelty to animals indicates low affective empathy and reduced inhibition to harm, dispositions that would importantly bear on their interactions with other humans.

Finally, it appears that people also categorise animals not just as non-human (and thus belonging to a distant outgroup), or as having lower mental capacities, but positively *as* animals. How can the Anti-Animal Bias be explained? People might be accustomed to treating animals as profoundly morally inferior and may associate being an animal with low moral status, whereas they are not accustomed to interacting with extraterrestrials. Moreover, a specific prejudice against animal-like beings could have been evolutionarily adaptive. Automatic classification of entities in one's environment as animals—as opposed to inanimate objects, plants or, of course, humans—was obviously evolutionarily adaptive. Realistic conflict theory assumes that conflict of interests between groups can lead to negative attitudes towards outgroup members (Sherif, 1961). Clearly, there are many conflicts of interests between humans and animals. Historically, animals either preyed on humans, or were preyed on by them (cf. Kasperbauer, 2017). Therefore, it is plausible that humans have developed a specific bias against animals that they might not have towards other kinds of non-humans.

Future Research

The psychological factors underlying moral anthropocentrism that I have identified need further study. Future research could investigate their relative strengths in contributing to moral anthropocentrism and uncover their underlying psychological mechanisms. In particular, future studies could investigate whether the factors I have identified are linked to behaviour in more realistic contexts, and whether they play a significant causal role in driving behaviour towards animals or rather reflect post-hoc justifications of such behaviour.

Future research could also investigate links between the factors of moral anthropocentrism and ‘folk biology’. For example, research on racism showed that perceiving racial group membership as determined by biological factors as opposed to by sociocultural factors (i.e. biological essentialism) influences people’s discriminatory views towards racial groups (Williams & Eberhardt, 2008). Similarly, it is possible that, for example, the Anti-Animal Bias is underpinned by an animalistic essentialist view. This raises the question of whether people apply the Anti-Animal Bias even to humans in cases of dehumanization of human outgroups.

In my studies the only animal species participants were presented with was monkeys. However, previous research has shown that people perceive some animals to be morally more important than others, e.g. dogs compared to pigs (cf. Chapter 2). Future research could investigate how the different factors underlying moral anthropocentrism play out when different animals are compared with each other, such as animals that are usually categorised as food, experimental subjects or pets. It is also possible that further factors underlying moral anthropocentrism will be discovered that go beyond the ones I have identified, e.g. biases against specific categories of animals, such as pests.

Conclusion

Humans prioritise themselves over all other animals. My aim was to identify the factors that underlie this attitude and present a first attempt of a framework for moral anthropocentrism,

which future research can study in-depth. I found that it is not restricted to animals, but applies to other kinds of hypothetical non-humans, and that it is not driven entirely, or even primarily, by perceived differences in mental capacity, whether at the individual or species level. Critically, it is underpinned not by a single factor but by several distinct ones. I found that people prioritise humans over animals because (1) animals have lower mental capacities (Individual Mental Capacity View); (2) people believe that individuals should prioritise members of their own species over others (Species-Relativism); (3) of a tendency to consider humans to be of superior value in an absolute sense (Pro-Human Bias); (4) of a specific bias against animals, a bias people do not manifest towards non-humans with equivalent mental capacities (Anti-Animal Bias). As I showed, these four factors can come apart in hypothetical scenarios. However, since all of them apply to all actual animals (perhaps with the exception of species with higher cognitive capacity, such as apes and dolphins), it is not surprising that moral anthropocentrism is such a robust and pervasive attitude. Attempts to show that we routinely underestimate the mental capacities of animals, or that we are linked with them by relations of similarity or fellowship, can influence only some aspects of moral anthropocentrism. Moreover, while differences in mental capacity, and partiality to group members, can arguably justify prioritizing humans in at least some contexts, Species-Absolutism—regarding humans as absolutely morally more important simply in virtue of being humans (Pro-Human Bias), and animals as absolutely morally inferior simply in virtue of being animals (Anti-Animal Bias)—clearly seems like a bias. My studies thus demonstrate that moral anthropocentrism is partly driven by prejudice. The interpretation that at least certain aspects of speciesism can be referred to as prejudice is in line with my findings of Chapter 2 that showed that speciesism correlates with SDO and human-human forms of prejudice.

4. Deontological Constraints Against Harming

Animals

In the previous two chapters, I demonstrated that speciesism exists as a psychological construct and explored its underlying factors. I found that, as predicted by the speciesism hypothesis, people value humans more than animals just due to their species-membership. In this chapter, I want to investigate a more specific implication of what it means to value humans more than animals. A key aspect of what it means to grant someone moral value is the question of whether and when we consider it permissible to harm that individual. If we consider an individual to have high moral status, this generally means that we consider it wrong to harm that individual, unless there are very good reasons to do so. Presumably, most people would consider it wrong to harm both humans and animals for *no* reason (e.g. people are against unnecessary deliberate animal cruelty). However, what if there *are* reasons to harm an individual, such as if that harm were the only way to prevent even more harm?

Utilitarianism tells us to always maximise the good. This is an attractive moral aim in many cases—few would deny that, for example, we should try to save as many lives as we can, and generally minimise harm to others. In some cases, however, most people think that it would be wrong to maximise the good because this would violate what moral philosophers call ‘deontological constraints’ (Nozick, 1974)—moral rules that forbid harming, torturing, or killing innocents, even for a good end. A great deal of recent psychological research has examined whether and under what circumstances people hold deontological constraints (e.g. Greene, 2014). For example, in the famous Footbridge trolley case (Thomson, 1984) a large majority of people are reluctant to push one person off a bridge to save the lives of five people trapped on the tracks of an oncoming train (Greene, 2009), and it has even been claimed that such constraints are universal and innate (Mikhail, 2007). What has not been studied yet is whether people also apply

deontological rules to animals. Are people as reluctant to harm animals to save a greater number of animals as they are for humans? And if indeed people do accept such constraints on harm to animals, are these deontological constraints similar in structure and strength to those that protect humans?

Kant (1785) offered one answer to these questions. He argued that the source of the deontological constraints that protect us from being sacrificed for the sake of others is our unique rational nature—our rationality makes us priceless and, therefore, we cannot be an input into a utilitarian cost-benefit calculation. On Kant’s own view, this meant that, because animals are not rational, they do not morally matter in their own right, though he frowned upon gratuitous cruelty to animals because it could make us more willing to harm humans. The philosopher Robert Nozick (1974) famously discussed an alternative principle he referred to as: “utilitarianism for animals, Kantianism for people”. On this view, animals do matter, but they matter very differently than humans. As non-rational individuals, they do not enjoy deontological protections and *can* be entered into utilitarian calculations.

Chapter 2 has demonstrated that people morally value humans more than animals primarily due to mere species-membership and to a smaller extent due to perceived level of sentience (i.e. capacity to suffer) or intelligence. Work on moral expansiveness shows that people represent different entities’ moral status along a moral circle of concern, where humans (in the centre) receive the highest and lower animals (outer boundaries) receive the lowest concern (see e.g. Laham, 2009; Crimston et al., 2016). Some of these studies also suggest that people perceive humans to be more deserving of prevention of harm than animals (e.g. Crimston et al., 2016). Further, in moral dilemmas in which the lives of humans are pitted against the lives of animals, people generally endorse harming animals to save humans (e.g. Topolski, Weaver, Martin, & McCoy, 2013; Petrinovich, O’neill, & Jorgensen, 1993). A recent large-scale survey assessing people’s judgments on how autonomous cars should act in moral dilemma situations showed that people exhibited a

very strong preference to save humans over animals (Awad et al., 2018; for a similar study, see Kim et al., 2018).

While it is clear that people value animals less than humans, it remains an open question how exactly they value them. In particular, I am interested in what types of deontological constraints against harming animals or humans people hold (e.g. in sacrificial moral dilemmas). Are people willing to harm a few animals or humans to save many animals or humans, and how do the respective deontological constraints differ from each other? It is obvious that people are willing to harm animals to save humans. Yet it is unclear, for example, to what extent people are willing to harm a few animals to save many animals, and how this tendency compares to their willingness to harm a few humans to save many humans. The answers to these questions are not obvious and multiple hypotheses seem plausible. From a certain perspective it could be argued, for example, that harming one to save many is a sign of greater moral concern because more beings will be saved. Given that people are less concerned about animals, one might expect people to be more willing to harm a few humans to save many humans than to harm a few animals to save many animals. Another possibility is that people's willingness to harm in the animal case is identical to their willingness to harm in the human case. After all, both the harmed beings as well as of the saved beings are granted proportionally lower moral status in the animal case. This means that the relative moral status of the harmed and helped beings should stay the same in both the human and animal case. Finally, it is also plausible that people are more willing to harm a few animals to save many animals than to harm a few humans to save many humans. Their aversion to harm animals may be reduced to a particularly strong extent. If that is the case, are deontological constraints against harming animals reduced only slightly or completely? Are people's intuitions about such dilemmas correctly captured by Kant, utilitarianism, Nozick, or by some other view? In this chapter I systematically test a range of possible descriptive ethical views of how deontological constraints against harm might apply to humans and animals.

Why is this of interest? First, answering this question can shed important light on the psychology of moral judgment more generally. Do deontological constraints pervade all moral thinking, or do they only apply in a distinctive domain? If deontological constraints apply only to humans, how does morality operate with regards to other species? Focusing on dilemmas where we can sacrifice some humans to save a greater number, Greene (2014) has suggested that people engage in either utilitarian or deontological modes of moral thinking, and that when the two compete, the intuitive, deontological mode typically prevails. If deontological constraints apply only to humans, this would mean that this dual process picture applies only in the human domain, and that moral judgment in the animal domain is the result of a single process. Conversely, if deontological constraints *do* apply to animals, this would raise the psychological question of whether these deontological constraints apply in the *same* way to animals as they do to humans. Moreover, since this would show that, as opposed to Kant, people generally do not base deontological constraints to rationality or other higher cognitive capacities, it would raise the further question of *what* is the basis of such constraints—is it mere sentience, or could it also be species-membership itself (i.e. speciesism)?

Second, how people think about deontological constraints against harming animals is also of intrinsic interest. Animals are routinely harmed to benefit humans (e.g. for consumption or for medical experiments), and in some cases animals are harmed to benefit *other* animals. But it is not yet clear whether most people nevertheless hold that there are moral constraints on when and whether such harm is permissible, and what psychological processes underlie such intuitions.

Relevant Previous Sacrificial Dilemma Research

While several previous studies have shown that people are willing to sacrifice animals to save humans, these studies did not investigate intra-species sacrifice cases, i.e. cases where one animal has to be harmed to save many animals. Therefore, these studies do not provide the evidence required to identify which ethical view of harm best captures lay people's intuitions.

A relevant recent study looked at choices pitting one animal versus other animals (Bostyn, Sevenhant, & Roets, 2018). In this study, participants were presented with two mice cages attached to an electric shock machine. They were informed that, as a default, the five mice in one of the cages would receive a painful electric shock unless participants decided to intervene and push a button that would instead redirect the electrical current to the other case that contains just one mouse. 84% of participants pushed the button. In contrast, they found that when the same situation was described hypothetically, only 66% said that they would push the button. This study suggests that people may be fairly willing to harm a few animals to save many animals. However, since the study did not directly compare deontological constraints against harming animals to deontological constraints against harming humans, we cannot draw a definite conclusion from it.

A recent study found that in Trolley-like scenarios people were more likely to harm robots than humans to save many humans (Nijssen, Müller, Baaren, & Paulus, 2019). The more human-like the robots were, the less likely people were to harm them, which was explained primarily due to the attribution of affective states and less due to the attribution of agency. Another study found that people are more likely to endorse harming stereotypical out-group members for the greater good than stereotypical in-group members (Cikara, Farnsworth, Harris, & Fiske, 2010). More specifically, participants were more likely to endorse harming a person with a low level of competence such as a homeless person than a person with a high level of competence such as a rich person. The person's level of warmth, however, did not affect people's deontological constraints. Finally, Nichols & Mallon, (2006) found that people apply deontological constraints even to inanimate objects, albeit very weakly: while people did not consider it overall wrong to destroy one object to prevent several objects from being destroyed in a Footbridge-like scenario, they did agree that a moral rule was broken in such a case. Overall, these studies suggest that attributes above and beyond species-membership, such as perceptions of an individual's mental capacities, can influence deontological constraints against harm.

When is Harm Towards Humans and Animals Permissible?

In order to investigate how people think about harm to animals vs. to humans, we first need to set out the theoretical options (see *Table 4.1*). My aim is to try to identify the view that best describes lay people’s intuitions. The first two views are well-known normative philosophical positions that I include for completeness sake even if they are unlikely to be psychologically accurate. The remaining views are all plausible descriptions of people’s intuitions and have sometimes even been described as such.

Table 4.1. Potential moral views of when it is permissible to harm humans and animals.

Harm:	Humans	Animals	Animals	Humans >
To save:	Humans	Animals	Humans	Animals
Utilitarianism	✓	✓	✓	✗
Kant’s view	✗	✓	✓	✓*
Cross Species Deontology	✗	✗	✗	✗
Nozick’s slogan	✗	✓	✓	✓
Multi-level Uniform Deontology	✗	✗	✓	✓
Multi-level Weighted Deontology	weighted	weighted	✓	✓

Note. Humans > Animals stands for humans have higher moral status than animals (even if they are equally sentient).

*While Kant thought that the moral status of animals is not fundamentally different from that of objects, he did think that harming animals gratuitously is wrong on instrumental grounds, because it can make us more willing to harm humans.

The first possibility is that people follow *utilitarianism* throughout. On this account, the same harm matters equally, regardless of who suffers it (weighted by their degree of sentience)—whether humans or animals (Bentham, 1780). This view is anti-speciesist, meaning that species-membership itself should not influence the moral status of an individual (Regan & Singer, 1989).

If people were strict utilitarians, they would consider it permissible (or even required) to sacrifice both humans and animals to promote the greater good (of both humans and animals). It is unlikely, however, that this view captures the intuition of most people, since, as described above, there is considerable evidence that people accept deontological constraints against harming humans for the greater good (e.g. Greene, 2014), and people tend to value animals less (Chapter 2).

Another possibility is that people follow *Kant's view*. On this account, only humans matter morally and therefore deserve deontological protection, whereas animals are just seen as objects that can be used to our own ends (Kant, 1785). The source of the value for humans is the fact that we, but not animals, are rational individuals and possess advanced cognitive capacities. This, however, is also implausible as an account of most people's view, since people do believe that animals matter morally at least to some extent (Chapter 3).

Yet another possibility is that people follow a view we can call *Cross Species Deontology*. On this account, the same deontological principles apply in the same way to all species. Neither humans nor animals should be sacrificed for the greater good of either. Like utilitarianism, this view is anti-speciesist because species-membership itself should make no difference (for a similar view, see Regan, 1987). Abolitionist animal rights activists tend to endorse similar moral positions (e.g. Francione, 1995). This again is unlikely to capture the common view given that people on average think that it is permissible to harm animals to benefit humans, e.g. via medical testing (Chapter 2).

There are, however, at least three ways to capture this intuitive moral difference between humans and animals while still ascribing some moral significance to animals, *contra* Kant. The first is the Nozick (1974) suggestion, discussed above, that deontology applies only to humans, while utilitarianism applies to animals. When it comes to animals, we should simply maximise utility—even by sacrificing some for the greater good. But when we turn to humans, this is forbidden. Many scholars take Nozick's suggestion very seriously (e.g. McMahan, 2002). Killoren and Streiffer

(2019) point out that many lay people's beliefs appear in line with this view. For example, people generally consider keeping animals in well-run zoos as morally permissible. Keeping innocent humans imprisoned, in contrast, is generally considered repugnant even if they were kept very happy. Another example is animal research. Most Western jurisdictions have a near absolute prohibition of more-than-minimal research on children who, like animals, cannot give consent (Ross 2005; Gennet & Altavilla, 2016)—even if it could be beneficial in expectation. Regulations and intuitions for research on animals are much more permissive (Varner, 1994). The Nuffield Council on Bioethics even called the utilitarian cost-benefit analysis the “cornerstone” of research on animals (Kilkenny, Browne, Cuthill, Emerson, & Altman, 2010; Nuffield Council on Bioethics, 2005).

A second possible view is that there is a *hierarchy* of moral status, reflected by species-membership in our case, i.e. speciesism. Individuals that are lower in the hierarchy (e.g. pigs) can be sacrificed for the sake of those higher up (e.g. humans). But *within* each level of moral status, the deontological constraints offer the same protections (i.e. it is wrong to sacrifice a pig to save five other pigs), and these protections apply to the *same* degree. Call this *Multi-level Uniform Deontology*.

A final, more complex view, which I call *Multi-level Weighted Deontology*, is that the deontological protections are not absolute, and get *weaker* the lower the level of moral status. As we go down the hierarchy, the less stringent the deontological constraints. According to this view, people would consider harming animals to save many animals neither completely permissible nor completely wrong (as Multi-level Uniform Deontology would) but instead somewhere in between. Further, the lower the moral status of the animal in question, the more permissible they would consider harming it to save many animals with the same moral status (i.e. it is more permissible to sacrifice one pig to save five pigs than to sacrifice one human to save five humans). In cases where the moral status of a being is very low, such as with animals, Multi-level Weighted Deontology

should resemble Nozick's slogan because in both cases the deontological constraints will be low or non-existent.

These three hypotheses have not been tested before. My own hypothesis is that Multi-level Weighted Deontology is the view that describes people's intuitions best. I expect people to consider it permissible to harm animals to benefit humans, and at the same time somewhat, but not completely, permissible to harm animals to save many animals with the same moral status. This is based on the notion that people do not regard deontological constraints as absolute (Kahane, 2015) but rather that constraints get weighted and can sometimes be overturned. For example, people likely do not consider it completely wrong to harm a few humans to save a *very* large number of humans. Similarly, I expect their constraints against harming animals to be neither absolute nor non-existent but rather moderate in strength (Holyoak & Powell, 2016).

The Present Research

In eight studies I tested the extent to which people accept deontological constraints against harming animals and compared it to the extent to which people accept such constraints for humans. In all studies participants were presented with different (usually hypothetical) sacrificial moral dilemmas, which involved harming either animals or humans in order to prevent harm to a larger number of animals or humans and asked for participants' judgments of the moral permissibility (rightness) of doing so. In Study 4.1, I tested whether people deem it more permissible to harm the few for the benefit of the many when both the victims and beneficiaries are animals, compared to when they are humans. In Study 4.2a and 4.2b, I explored whether this effect varies depending on the animal species (pigs, dogs or chimpanzees). In Study 4.3, I investigated how moral judgments change when the victims and beneficiaries belong to different species. In Study 4.4, I tried to experimentally manipulate speciesism to observe a change in the strength of the deontological constraint. In Study 4.5, I contrasted the extent to which people hold deontological constraints for humans, animals and mere objects. In Study 4.6, I tested to what

extent an individual's cognitive capacity level and sentience contribute to the degree of deontological constraint applied to the individual. Finally, in Study 4.7, I replicated my main effect in a different population and linked judgements to behaviour in the real-life context of medical experimentation.

Open Science. As for all studies in this chapter, reports of all measures, manipulations, and exclusions, and all data, analysis code, and experimental materials are available for download at: https://osf.io/nt69s/?view_only=ceb54033756b42c38397676830497b8e. Studies 4.2a, 4.2b, 4.5, and 4.6 were pre-registered on the Open Science Framework.

Ethics Statement. For all studies, relevant ethical guidelines were followed, and the research was approved through University of Oxford's Central University Research Ethics Committee, with the reference numbers MS-IDREC-R56657/RE001 and MS-IDREC-R58592/RE001.

Study 4.1. Humans vs. Animals

In the first study I aimed to investigate whether people's deontological constraint against harming humans differs from their deontological constraint against harming animals. In particular, I tested whether people consider it more permissible to harm one pig in order to save five pigs than to harm one human in order to save five humans. My hypothesis was that this was the case, but that people would still not consider it completely permissible to harm the pig. My reasoning was based on the Multi-level Weighted Deontology hypothesis, namely that people's deontological constraints are weighted by the moral status they attribute to the respective individual. Since people attribute much higher moral status to humans than to pigs, their deontological constraints will be much higher for humans than for pigs. Relatedly, I hypothesised that individual differences in the level of speciesism—as measured by the Speciesism Scale (Chapter 2)—would correlate with moral judgments about harm to pigs (but not harm to humans), i.e. the extent to which people hold

deontological constraints for pigs. This is because the more speciesist people are, the lower the moral status they attribute to pigs, and as a consequence, one would expect the deontological constraints to be weaker.

A secondary hypothesis was that the difference in deontological strengths between humans and pigs would be stronger in personal sacrificial moral dilemmas such as the Footbridge case than in impersonal sacrificial moral dilemmas such as the Switch case because harm aversion will be more prominent in the former than the latter case (Greene, 2014). Finally, I hypothesised that individual differences in the endorsement of *instrumental harm* (OUS-IH; Kahane et al., 2018), but not of *impartial beneficence* (OUS-IB) would correlate positively with moral judgments.

Method

Participants. 332 US American participants took part in the study online via MTurk and received \$0.60 payment (in line with US minimum wage for all studies) for their participation. Sixteen participants were excluded for either failing to complete the questionnaire or for failing an attention check, leaving a final sample of 316 people (134 female; $M_{age} = 36.49$, $SD = 11.62$). I aimed for a sample size of 300 participants, as my a priori power analysis showed that 301 participants were required to detect a medium effect size of $f = 0.18$, with an alpha of 0.05, and power of 0.80 (Cohen et al., 2014).

Design, Materials and Procedure. My study had a 2 dilemma (impersonal vs. personal) x 2 species (humans vs. pigs) between-subjects factorial design. Participants were randomly allocated to read one of four variants of a moral dilemma modelled after the traditional trolley problem. As a personal dilemma the Footbridge trolley case was used and as an impersonal dilemma the Switch trolley case was used (see online materials). Participants were asked to indicate

how morally right (i.e. permissible⁵) or wrong they thought it to be to harm one pig (human) to save five pigs (humans) on a 7-point scale from “absolutely morally wrong” to “absolutely morally required” (from now on referred to as moral judgment throughout the chapter). Then, participants completed the Speciesism Scale (Chapter 2), consisting of 6 items such as “Morally, animals always count for less than humans”, which captures individual differences in speciesist attitudes (internal consistency: $\alpha = .88$). Next, participants completed the Oxford Utilitarianism Scale (Kahane et al., 2018) which measures general proto-utilitarian tendencies. The *instrumental harm* subscale (OUS-IH) measures a ‘negative’, permissive attitude toward instrumental harm, consisting of 4 items such as “Sometimes it is morally necessary for innocent people to die as collateral damage—if more people are saved overall” (internal consistency: $\alpha = .77$). The *impartial beneficence* subscale (OUS-IB) measures a ‘positive’, impartial concern for the greater good, consisting of 5 items such as “It is just as wrong to fail to help someone as it is to actively harm them yourself” (internal consistency: $\alpha = .72$). Finally, participants responded to demographic items.

Results

A 2x2 between-subjects ANOVA revealed two main effects but no interaction. Participants considered harming one pig to save five pigs to be more permissible (impersonal: $M = 5.10$, $SD = 1.44$; personal: $M = 4.23$, $SD = 1.81$) than harming one human to save five humans (impersonal: $M = 4.21$, $SD = 1.59$; personal: $M = 3.04$, $SD = 1.96$), $F(1, 312) = 28.95$, $p < .001$, $\eta_p^2 = .09$. Participants also considered harming in an impersonal dilemma (i.e. via flipping a switch) to be more permissible than in a personal dilemma (i.e. via pushing off the footbridge), $F(1, 312) = 28.08$, $p < .001$, $\eta_p^2 = .08$. There was no interaction in moral judgments between species and dilemma, $F(1, 312) = 0.62$, $p = .43$, $\eta_p^2 = .002$ (Figure 4.1). On a descriptive level, the personal-

⁵ For reasons of simplicity and philosophical accuracy I use the term “permissible” even if the study question included the terms “required” or “right”. While these are not fully identical on a conceptual level, Study 3 revealed that participants do not clearly distinguish between them in our dilemmas.

humans condition seemed most polarizing, with a sizable portion of participants being strictly against harming humans for the greater good (cf. distributions of answers in *Figure 4.1*).

The more speciesist participants were, as measured by their scores on the Speciesism Scale, the more permissible they considered it to harm one pig to save five pigs, $r = .26$, $p < .001$ across both pig conditions. In contrast, speciesism did not significantly correlate with moral judgments in the human conditions, $r = .12$, $p = .13$. The two correlation coefficients did not differ significantly from each other, $z = 1.28$, $p = .20$. The more participants endorsed instrumental harm, as measured by OUS, the more permissible they considered harming one for many in all four conditions, with an overall correlation of $r = .42$, $p < .001$. Endorsement of impartial beneficence did not correlate with moral judgments. There were no noteworthy correlations between moral judgment and demographics. It is worth noting that the more speciesist people were, the less they endorsed impartial beneficence ($r = -.19$, $p < .001$) and the more they endorsed instrumental harm ($r = .32$, $p < .001$).

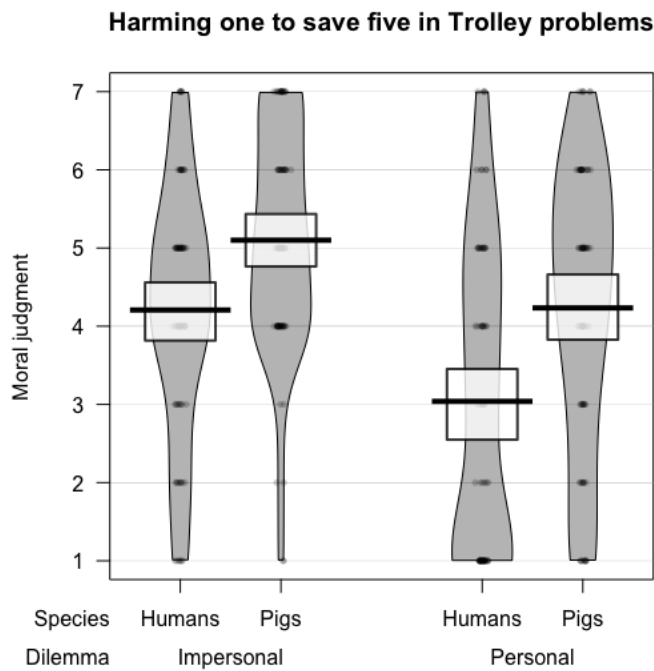


Figure 4.1. Moral judgments on harming one to save five individuals, ranging from 1 (*morally wrong*), over 4 (*neither right nor wrong*), to 7 (*morally right*). Black points represent raw data, horizontal bars represent means, rectangles represent confidence intervals, and “beans” represent smoothed densities.

Discussion

Results from Study 4.1 confirm my main hypothesis that people consider it more permissible to harm one pig to save five pigs than to harm one human to save five humans. Contrary to my secondary hypothesis, the effect occurred equally in both personal and impersonal dilemmas, which suggests that harming humans vs. animals is not analogous to personal vs. impersonal sacrificial dilemmas. Overall, these findings suggest that people have weaker deontological constraints for animals than for humans and are therefore more likely to make a utilitarian cost-benefit calculation when confronted with animals than with humans. It was, however, not the case that people considered it completely permissible to harm animals to save

many, which is what Nozick's slogan "utilitarianism for animals, Kantianism for people" would suggest, nor was the strength of the constraint constant across the case. The results therefore support the Multi-level Weighted Deontology hypothesis on which deontological constraints still limit the permissibility of harm to animals, but in a weaker way than in the human case. Note that Multi-level Weighted Deontology does not make predictions regarding the absolute level of permissibility, e.g. whether people tend to consider harming pigs permissible or impermissible (i.e. above or below the mid-point). Rather, the point is that constraints continuously increase or decrease.

While this is a description of the average response, a look at *Figure 4.1* makes clear that, while there were large individual differences in responses, only a small minority held *no* deontological constraint against harm for animals, which is what Nozick's slogan would predict. One could argue, however, that even if Nozick's slogan was broadly true, due to individual variance and noise it is unrealistic to expect the average response to be at 1, i.e. indicating that on average people have zero deontological constraints for animals. Therefore, a clearer way to disentangle Nozick's slogan from Multi-level Weighted Deontology is to test whether the strength of deontological constraints for different types of animals, which people value to different degrees, differs slightly. If they do, that would give clearer support for the Multi-level Weighted Deontology model and against Nozick's slogan.

Study 4.2a. Different Types of Animals

In the second study I aimed to replicate my finding from Study 4.1 that people hold weaker deontological constraints against harming animals than against harming humans. Furthermore, I aimed to extend these findings to other types of animals: chimpanzees and dogs. This will help clarify whether the extent to which people value different animals influences the strength of the deontological constraint as predicted by Multi-level Weighted Deontology (but not by Nozick's

slogan). Dogs and chimpanzees are interesting cases because they both typically are valued less than humans and more than pigs despite the fact that dogs and pigs are similar to each other in terms of level of intelligence or emotional capabilities (Chapter 2). Chimpanzees are the animal species that are most closely related to humans but still receive a much lower moral status than humans (Chapter 2). For these reasons, it is of particular interest to compare the deontological constraints people hold for these four types of species.

My hypothesis, which I pre-registered at <https://osf.io/axq9z/>, was that, as in Study 4.1, people would consider it more permissible to harm one pig to save many pigs than to harm one human to save many humans. Furthermore, I hypothesised that moral judgments about harming chimpanzees or dogs would lie between moral judgments about pigs and humans. I did not have a hypothesis whether moral judgments about harming chimpanzees would differ from moral judgments about harming dogs.

Method

Participants. I recruited 201 US American participants online via MTurk who received \$0.60 payment (in line with US minimum wage) for their participation. Seven participants were excluded for failing an attention check or not completing the study, leaving a final sample of 194 people (97 female; $M_{age} = 34.18$, $SD = 10.53$). As the effect size between humans and pigs in Study 4.1 was $f = 0.31$, I assumed $f = .25$ for this study because I expected a weaker effect size for the new dogs/chimpanzees vs. humans comparisons. My analysis showed that 180 participants were required to detect this effect with an alpha of 0.05, and power of 0.80 (Cohen et al., 2014; Soper, 2017), so I aimed to recruit 200 participants to account for any exclusions.

Design, Materials and Procedure. This study had four between-subjects conditions: pigs, dogs, chimpanzees, and humans. As in Study 4.1, participants were randomly allocated to one of four variants of an abstract sacrificial dilemma scenario:

*“Suppose you are in a situation in which you have to decide whether to kill one **pig** to save five **pigs**. If you do nothing, the five **pigs** will die. The only way to save them is to kill the one **pig**.”*

In the respective other conditions, the word pig(s) was replaced with dog(s), chimpanzee(s), or human(s). After participants read the dilemma, as in Study 4.1, they indicated how morally right or wrong it would be to sacrifice one to save five on a 7-point scale. Participants then completed the Speciesism Scale and Oxford Utilitarianism Scale. Next, participants rated each of the four species on a scale from 0-100 according to the following measure: intelligence, moral value, capacity for pain, capacity for fear, and how much empathy participants feel towards individuals of this species. These items were included for exploratory reasons. Finally, participants responded to demographic items.

Results

An ANOVA across the whole design revealed that participants' moral judgments were significantly different across experimental conditions, $F(3, 190) = 2.91, p = .036, \eta^2 = .03$ (Figure 4.2). Post-hoc tests showed that participants judged instrumentally harming humans to save other humans as more morally wrong ($M = 4.12, SD = 2.12$) than doing so for pigs ($M = 5.08, SD = 1.46; p = .026$; Tukey's HSD). There were no other significant differences between species. Moral judgments between the chimpanzees ($M = 4.84, SD = 1.42$) and dogs ($M = 4.77, SD = 1.58$) conditions did not significantly differ from each other ($p = .99$), and neither did they differ from moral judgments in the humans ($p = .15$ for chimpanzees; $p = .24$ for dogs) or pig ($p = .33$ for chimpanzees; $p = .34$ for dogs) condition.

As in Study 4.1, the more speciesist participants were, the more moral they considered instrumentally harming one animal to save five animals (all animal conditions collapsed: $r = .21, p = .01$). Yet, speciesism did not affect how people judged instrumentally harming one human to save many humans ($r = .12, p = .42$). The two correlation coefficients did not significantly differ from each other, $z = 0.56, p = 0.29$. Instrumental harm (OUS) correlated significantly with moral

judgments across all four conditions ($r = .44, p < .001$). Impartial beneficence (OUS) did not correlate with moral judgments ($r = -.02, p = .74$).

Descriptively, the variance in moral judgments in the human condition was higher than in all three animal conditions (*Figure 4.2*). As in the personal dilemma in Study 4.1, the human condition was more polarizing than the three animal conditions. A large proportion of participants believed it is absolutely morally wrong to sacrifice one human to save many humans, whereas this is not the case for the animal conditions.

Participants on the 100-points scale rated the moral value of humans the highest ($M = 91.91, SD = 14.246$), followed by chimpanzees ($M = 70.84, SD = 27.06$), followed by dogs ($M = 69.09, SD = 28.245$), followed by pigs ($M = 54.81, SD = 32.57$). The moral value attributed to pigs was significantly lower than the moral value attributed to dogs, $t(193) = 9.52, p < .001$. This ranking is in line with the strengths of the deontological constraints we observed people hold for these different types of species. There were no relevant correlations between moral judgements and the exploratory measures I included, which might be due to the small sample size in each condition. And there were also no noteworthy correlations with demographic variables.

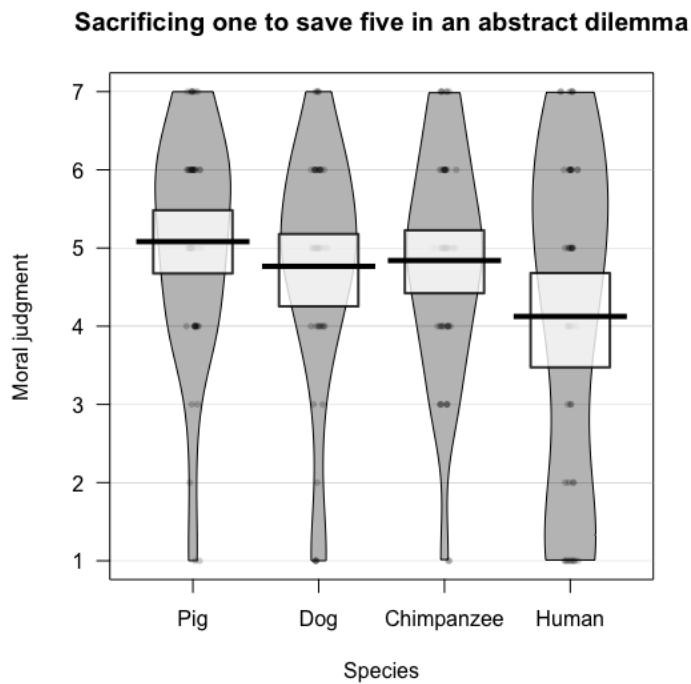


Figure 4.2. Moral judgments ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*).

Discussion

I replicated my finding from Study 4.1 that people consider it more moral to instrumentally harm one pig to save many pigs than to instrumentally harm one human to save many humans. Again, I found that the more speciesists people were, more willing they were to instrumentally harm animals of any type. The descriptive values for the chimpanzees and dogs conditions lie right in between the ones in the pigs and humans conditions, but did not differ statistically significantly from neither. Therefore, I was unable to draw a conclusion whether these two species are more judged like humans or more judged like pigs. I furthermore found that people generally seem to hold stronger deontological constraints for beings the higher their perceived moral standing is. People explicitly attribute higher moral standing to humans than to all types of animals, which might be taken as a further indication that the moral standing of individuals has a direct influence on the extent of the deontological rule people hold for them.

In hindsight I suspected that Study 4.2a was underpowered. My power analysis relied on an effect size of $f = .25$, which I chose because I expected a slightly weaker effect for the new dogs/chimpanzees vs. humans comparisons. However, looking at the results, I have overestimated the effect and instead should have chosen a smaller expected effect size of $f = 0.13$. I address these issues in Study 4.2b.

Study 4.2b. Different Types of Animals

In Studies 4.1 and 4.2a I found that participants judged it to be significantly more moral to instrumentally harm one pig to save several pigs than to instrumentally harm one human to save several humans. However, two further experimental conditions in Study 2a using chimpanzees and dogs, respectively, did neither significantly differ from the pig nor from the human condition, which is why no definitive conclusion can be drawn from these results.

The purpose of Study 4.2b, which I pre-registered at <https://osf.io/8qynr/>, was to again explore whether people consider instrumentally harming one chimpanzee to save many or one dog to save many to be more or less morally right than instrumentally harming one human or one pig to save many. As such, this study can be seen as a partial replication of Study 4.2a, with two important differences. Most crucially, I relied on a higher sample size. Further, I developed and pre-tested (see below) a new scenario that was seen to be of higher realism both than the abstract one of Study 4.2a and the standard trolley paradigm, which also seems detached from real-life cases. This way, I believed to address the shortcomings of Study 4.2a.

My hypothesis was that moral judgments of instrumentally harming a few beings to save several more beings differ between humans and animals, as well as between different animals. I did not have a hypothesis regarding whether chimpanzees or dogs would be closer to pigs or humans. Again, I predicted that the higher people score on speciesism, the more moral they judge sacrificing animals to save several more animals.

Method

Development and Pre-test of Materials. One criticism of the traditional Trolley problem is that people consider it unrealistic and thus do not take it morally serious enough (Bauman, McGraw, Bartels, & Warren, 2014), an issue that is exacerbated when the scenarios involve animals. Therefore, in a pre-test to Study 4.2 ($N = 700$; 33-37 participants per group) I used five different sacrificial dilemmas and compared them to each other: the former two I devised myself, the latter three were adapted from the literature (Greene, Morelli, Lowenberg, Nystrom, & Cohen, 2008). The Vaccine Death dilemma involved harming ten individuals by developing a vaccine to save the lives of a hundred infected individuals. The Vaccine Pain dilemma was identical but instead of killing individuals to save others' lives, it involved inflicting strong pain on some to prevent strong pain in others. In the Transplant dilemma, one individual had to be killed in order to take his organs to save the lives of five others with different organ failures. Finally, the standard Footbridge and Switch Trolley dilemmas were included. Participants rated how silly (i.e. how realistic) they found each scenario. My aim was to find the scenario that is generally perceived to be most realistic irrespective of the type of species involved. A 5 (dilemmas) x 4 (species) between-subjects ANOVA revealed that the dilemmas were not perceived as equally realistic, $F(1, 688) = 30.04, p < .001, \eta^2 = .04$. The dilemma participants judged to be most realistic was the Vaccine Death dilemma ($M = 2.53; SD = 1.72$; on a scale from 1 *Not silly* at all to 7 *Extremely silly*). A one-way ANOVA found that within this scenario there was no difference in perceived realism between the four species $F(1, 138) = 2.72, p = .10, \eta^2 = .02$.

Therefore, I have decided to rely on the Vaccine Death dilemma in the next studies, which reads as follows in the pigs condition (where the words 'pig' and 'pigs' are replaced accordingly in the other conditions) and involves a shift in numbers of harmed and saved individuals compared to the traditional Trolley problem:

*“There is a sudden outbreak of a rare virus in a small town. It infects the local **pigs** and is fatal. If nothing is done, about 100 **pigs** are expected to die. You are a researcher one town over, and your lab is attached to the local veterinary clinic. The only way to save the 100 **pigs** is for you to infect 10 healthy **pigs** in the clinic to identify the vaccine that can then be used to prevent the virus from killing the other 100 **pigs**. However, this will kill the 10 **pigs** you infected that otherwise would have lived. Doing so will be straightforward, safe for you, and no one will find out.”*

Participants. I recruited 660 US American participants online via MTurk who received \$0.50 payment for their participation. Twenty-two participants were excluded for failing at least one of two attention checks, leaving a final sample of 638 people (292 female; $M_{age} = 35.69$, $SD = 10.49$). I anticipated a small to medium effect of $f = 0.13$ between animal species based on the results the previous study that relied on abstract dilemmas mentioned above. With an alpha of 0.05, and power of 0.80, my a priori power analysis showed that I required 652 participants (Cohen et al., 2014; Soper, 2017). Hence, I aimed to recruit 660 participants to account for exclusions.

Design, Materials, and Procedure. This study had four between-subjects conditions to which participants were randomly allocated to: pigs, chimpanzees, dogs, or humans. After participants read the vaccine dilemma, they indicated how morally right or wrong it would be to kill ten to save a hundred on a 7-point scale as well as on a binary choice variable (right or wrong).

For exploratory purposes, participants also answered how many individuals would need to be saved (minimum) for it to be morally right to kill ten. Participants could either enter a number or select “it is never right”. And participants indicated how morally acceptable (7-point scale) and how morally required (7-point scale) they considered it to kill 10 to save 100.

Participants then completed the Speciesism Scale and the Oxford Utilitarianism Scale. Since speciesism correlates with empathic concern (Chapter 2) and empathic concern correlates with instrumental harm (Kahane et al., 2018) I included the Empathic Concern (EC) Scale as a covariate. In particular, I wanted to test whether speciesism still correlates with moral judgments

in the animal conditions, but not the human condition, if empathic concern is controlled for. The EC scale forms part of the Interpersonal Reactivity Index (IRI; Davis, 1980) and consists of 7 items such as “When I see someone being taken advantage of, I feel kind of protective toward them” ($\alpha = .91$). Finally, participants responded to demographic items.

Results

A one-way ANOVA revealed that, as predicted, moral judgment differed significantly across species, $F(3, 634) = 35.32, p < .001, \eta^2 = .14$. Participants judged harming humans to save other humans ($M = 2.65, SD = 1.90$) as more morally wrong than doing so for pigs ($M = 4.61, SD = 1.93; p < .001$; Tukey’s HSD), dogs ($M = 3.99, SD = 1.84; p < .001$), and chimpanzees ($M = 4.36, SD = 1.69; p < .001$). Participants also judged harming dogs as more morally wrong than pigs ($p = .01$), but not chimps ($p = .28$; see *Table 4.2* below for M and SD). There was no significant difference in moral judgment of harming pigs compared to chimpanzees ($p = .22$). Similar patterns were found for the binary question of whether harming ten to save a hundred was considered morally right or morally wrong. 25.6% considered it right for humans, 58.4% for chimpanzees, 52.5% for dogs, and, 73.3% considered it right for pigs. Logistic regression revealed significant differences across conditions in this binary response variable, $\chi^2(1) = 60.92, p < .001, R^2 = .12$.

Again, participants judged harming a few animals to save many animals more permissible the more speciesist they were, yet this association was absent in the human condition (*Table 4.2*). This association held even after controlling for empathic concern (EC), which is itself negatively associated with speciesism (Chapter 2). Correlations between empathic concern and moral judgments were only significant in the pig condition ($r = -.16, p = .03$). The correlation coefficients between speciesism and endorsement of harming humans was significantly different to the correlation coefficient between speciesism and endorsement of harming animals (all animal conditions collapsed), $z = 4.39, p < .001$. Furthermore, results of a linear regression indicated that there was a significant interaction effect ($b = -0.13, t = -4.46, p < .001$) between the condition (i.e.

the species-membership) and individual differences in level of speciesism (i.e. the Speciesism Scale) on moral judgments, $F(7, 630) = 23.2, p < .001, R^2 = .21$. The regression model that included the interaction term was significantly better at explaining the data than a model without the interaction term and just additive effects, $F(3) = 6.87, p < .001$. This supports the view that speciesism reduces deontological constraints against harming animals but does not affect deontological constraints against harming humans. And as in the previous studies, participants judged harming individuals for the greater good as more permissible, the more strongly they endorsed instrumental harm (OUS-IH) in all conditions (Table 4.2).

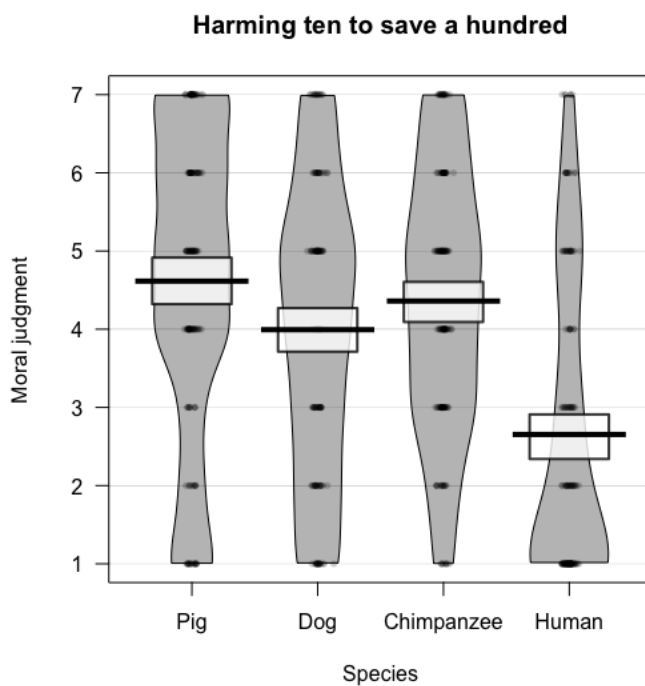


Figure 4.3. Moral judgments of harming 10 individuals of a species to saving 100 of the same species, ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*).

Table 4.2. Moral judgments and correlations (Study 4.2)

	Never right	Tipping Point M (Median)	Correlations			
			Speciesism	Speciesism (EC controlled)	OUS-IH	OUS-IB
Pig	25.5%	931 (78)	.25**	.20**	.37**	-.14
Dog	33.8%	1,615 (100)	.36**	.34**	.42**	-.04
Chimpanzee	29.2%	432 (99)	.32**	.30**	.46**	.17*
Human	59.6%	26,446 (100)	-.08	-.13	.50**	.31**

Note. * $p < .05$. ** $p < .01$. *Never right* stands for the proportion of participants who thought that it was never right to harm individuals of a certain species irrespective of the number of saved individuals of the same species. *Tipping point* stands for the number of individuals that have to be saved in order to justify harming ten individuals of the same species. *EC* stands for Empathic Concern.

As can be seen in Table 4.2, participants were much more likely to state that it is never right to harm ten humans than to harm any type of animals, no matter what number of individuals can be saved. Those who believed that there was a number of saved individuals that would justify harming ten individuals (Tipping point) stated much larger numbers for humans than for the three types of animals. I found no noteworthy differences in patterns between moral judgments and the exploratory measures of moral rightness and moral acceptability.

Participant age correlated with moral judgements such that younger participants were more likely to judge harming humans as more permissible, $r = -.31$, $p < .001$, but not animals, $r = -.03$, $p = .54$. Furthermore, men were more likely to judge harming individuals for the greater good to be permissible than women, $r = -.12$, $p < .001$. Apart from that there were no noteworthy correlations with demographic variables.

Discussion

The results of Study 4.2b demonstrate that people consider it significantly more wrong to harm a few humans to save many than to harm a few animals to save many, irrespective of the type of animal—whether they are pigs, chimpanzees, or dogs. In Study 4.2a I did not find conclusive evidence for whether this is also the case with chimpanzees and dogs, possibly due to a too small sample size and due to the use of a scenario that lacked psychological realism. In this study I changed both of these aspects. A number of further measures I included all broadly converged and support these main findings. For example, about 60% of participants believed that it is never permissible to harm ten humans no matter how many can be saved, whereas only ca. 30% believed this was the case for harming animals. I also found that people considered it more wrong to harm dogs or chimpanzees than pigs or chimpanzees to save many of the same species. Depending on the measure, chimpanzees were sometimes closer to pigs and sometimes closer to dogs. But in all measures, people clearly considered it most morally acceptable to harm pigs, and least acceptable to harm humans, whereas dogs and chimpanzees were in between.

I also found that the correlations between speciesism and moral judgments in the three animal conditions (and the non-correlation in the human condition) held even after controlling for empathic concern. This suggests that general empathic concern alone cannot explain people's weaker deontological constraints for animals than humans, but that instead speciesism itself is likely to be a strong driver of it. This was supported by my finding that species-membership moderated the effect of speciesism on moral judgments. While speciesism did affect moral judgments when animals had to be harmed, that was not the case when humans had to be harmed.

I observed that people have the weakest deontological constraints for pigs, somewhat weak ones for dogs and chimpanzees, and the strongest ones for humans. This supports the view that, indeed, deontological constraint against harming an individual to save more individuals of the same species depends on how much people value individuals of the respective species. In Study 4.2a I

found that participants rated the moral value of these four species in the following order: humans were highest, followed by chimpanzees and dogs, followed by pigs. This ranking is in line with the strengths of the deontological constraints we observed people hold for these different types of species. Crucially, the fact that deontological constraints differ slightly across different animal species rules out Nozick's slogan, since he did not predict that deontological constraints would differ across different animal species. Instead, these findings fit well with the Multi-level Weighted Deontology view, according to which deontological constraints get weighted differently depending on how much people value the respective species.

Study 4.3. Cross-Species Sacrifices

In the first studies, I found that people considered it more permissible to harm a few animals to save many animals than to harm a few humans to save many humans. In all conditions of the previous studies the individuals that could be saved were of the same species as those that had to be harmed. This means that the relative moral status of the individuals that have to be harmed and those that can be saved was kept constant. This was important to establish my main effect that people are more likely to endorse harming animals than humans if the relative value of the costs and benefits is held constant. In this study I aim to gain a deeper understanding of the mechanism of this effect by manipulating the relative value of the costs and benefits. I do this by examining cross-species sacrifices, i.e. cases when the species-membership of the harmed and the saved individuals differs.

According to the Multi-level Weighted Deontology model, deontological constraints are weaker the lower the moral status of the individual that has to be harmed. Further, the weaker the deontological constraints, the stronger the weight of the utilitarian cost-benefit analysis. From this, I draw the following hypotheses.

First, I would expect people to generally consider it more permissible to harm animals than humans irrespective of whether many animals or humans can be saved. This is because in both cases deontological constraints against harming animals are weaker, which makes it more likely for the utilitarian cost-benefit analysis to overrule them. However, since humans have higher moral status than animals, deontological constraints might be even easier overturned if humans compared to animals can be saved.

Second, since deontological constraints against harming humans are strong, people will in general consider it fairly wrong to harm humans to save either humans or animals. However, to the extent that people overcome their deontological constraints and do engage in cost-benefit analyses, these will yield opposing recommendations. Cost-benefit analysis will recommend harming the human to save many humans but advise against harming a human to save many animals. This is because most people consider ten human deaths much worse than one hundred pig deaths. I therefore expect an interaction effect such that the difference in moral judgments is greater if humans have to be harmed than if animals have to be harmed.

Method

Participants. 150 US American participants took part in the study online via MTurk and received \$0.90 payment for their participation. Eight participants were excluded for failing at least one of two attention checks, leaving a final sample of 142 people (66 female; $M_{age} = 38.60$, $SD = 11.66$). I anticipated a small to medium effect of $f = 0.13$ between animal species based on the results of Study 4.2. With an alpha of 0.05, and power of 0.80, my a priori power analysis showed that I required 120 participants for a within-subjects study (Cohen et al., 2014; Soper, 2017). To account for exclusions, I recruited 150 participants.

Design, Materials and Procedure. In this study I orthogonally manipulated both the species-membership of the individuals that can be saved and the species of individuals that have to be harmed. As such, this study had a 2 harm (pigs vs. humans) x 2 save (pigs vs. humans) design.

Note that in contrast to the previous studies, this study was conducted completely within-subjects to test whether the effect also holds in this set-up. In all other studies of this chapter, I relied on between-subjects designs in order to avoid potential carry-over effects.

I chose an abstract sacrificial dilemma because its brevity was well suited for this within-subjects design and because I expected large effects. Participants were asked to consider “Suppose you are in a situation in which you have to decide whether to kill 10 individuals to save 100 individuals. If you do nothing, the 100 individuals will die. How morally right or wrong is it to... 1) kill 10 humans to save 100 pigs, 2) kill 10 humans to save 100 humans, 3) kill 10 pigs to save 100 pigs, 4) kill 10 pigs to save 100 humans?” Participants responded to these four questions first on a 7-point moral judgment rating scale ranging from “*Absolutely morally wrong*” to “*Absolutely morally right*”, and on the next page on a binary response measure (*Wrong / Right*). On the next pages, I included two new questions on moral acceptability and more requiredness for all four conditions each to test whether they would yield different patterns. Next, I again included an exploratory question, asking the minimum number of individuals that would need to be saved in order to justify harming ten. Furthermore, I included a question asking the maximum number of individuals that could be harmed in order for it still to be justified to save the hundred. As in the previous studies, participants then completed the Speciesism Scale, the Oxford Utilitarianism Scale, and responded to demographic items.

Results

A two-way repeated-measures ANOVA revealed two main effects and an interaction (Figure 4.4). Participants generally considered it more permissible to harm pigs than humans, $F(1, 564) = 339.15, p < .001, \eta_p^2 = .38$, and generally considered it more permissible to harm an individual if humans could be saved than if pigs could be saved $F(1, 564) = 131.94, p < .001, \eta_p^2 = .19$. When humans had to be harmed, participants considered this much more permissible if this had to be done in order to save other humans ($M = 4.12, SD = 1.70$) than if it had to be done in

order to save pigs ($M = 1.83$, $SD = 1.36$). But, in contrast, when pigs had to be harmed, participants considered this only very slightly more permissible if this had to be done in order to save humans ($M = 5.78$, $SD = 1.54$) than if it had to be done in order to save pigs ($M = 5.04$, $SD = 1.68$). In short, there was a significant interaction in moral judgment based on species harmed and species saved (*Table 4.3*), $F(1, 564) = 34.07$, $p < .001$, $\eta_p^2 = .06$.

The same pattern was found for the binary morally right/wrong question: Only 4.2% considered it right to harm humans to save many pigs, but 51.4% considered it right to harm humans to save many humans. On the other hand, 80.3% considered it right to harm pigs to save many pigs, and 85.9% considered it right to harm pigs to save many humans. Results for the moral acceptability and moral requiredness questions were very similar to the main dependent variable (*morally right/wrong*; *Table 4.4*).

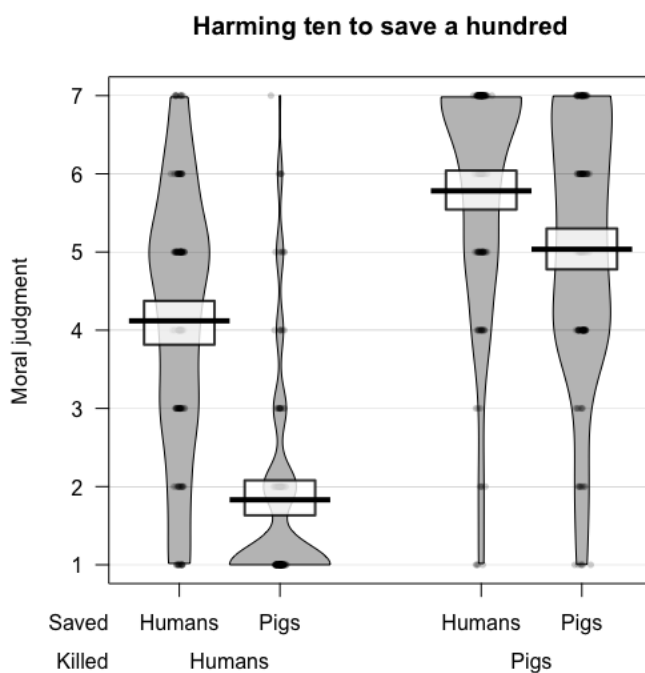


Figure 4.4. Moral judgments of killing 10 individuals of a species to saving 100 of either the same or a different species, ranging from 1 (*absolutely morally wrong*), to 4 (*neither right nor wrong*), to 7 (*absolutely morally right*).

Participants judged harming pigs, but not humans, more permissible the more speciesist they were (*Table 4.3*). However, speciesism did not correlate more with moral judgments depending on whether pigs or humans could be saved. A linear regression confirmed this: the saved-species factor ($b = 0.46, p = .02$), the killed-species factor ($b = -0.41, p = .03$), speciesism ($b = 0.17, p < .001$), as well as the interaction between speciesism and killed-species ($b = -0.22, p < .001$) were significant predictors of moral judgments, but not the interaction term between speciesism and saved-species ($b = 0.08, p = .1$).

Many more participants said it would always be morally wrong, no matter the number of saved individuals, to harm humans than pigs (*Table 4.4*). The type of saved species had a stronger influence on this when humans had to be harmed than pigs. Those participants who said that there was a number of saved individuals that would justify harming a few stated much higher numbers if humans had to be harmed than pigs. Similarly, more participants said it is always right to harm any amount of pigs than humans to save 100 individuals. Those participants who said that it depends on the number, stated much lower maximum numbers of harmed humans than pigs that would justify harming them to save 100 individuals. Furthermore, they judged harming individuals more permissible the more strongly they endorsed instrumental harm (OUS-IH) across all conditions, except in the sacrifice humans-save pigs condition (*Table 4.3*). There were no noteworthy correlations with demographics.

Table 4.3. Moral Judgments and Correlations (Study 4.3)

Species Killed	Species Saved	Morally wrong/right	Morally acceptable	Morally required	Correlations with Moral Judgment	
		<i>M (SD)</i>	<i>M (SD)</i>	<i>M (SD)</i>	Speciesism	OUS-IH
Human	Human	4.12 (1.70)	3.92 (1.76)	3.49 (2.04)	.04	.42**
	Pig	1.83 (1.36)	1.61 (1.19)	1.44 (1.09)	-.15	.07
Pig	Human	5.78 (1.54)	5.81 (1.65)	5.37 (2.08)	.39**	-.30
	Pig	5.04 (1.68)	5.28 (1.64)	4.27 (2.10)	.26**	.25**

Note. * $p < .05$. ** $p < .01$. The scales of the three measures of moral judgement range from 1 (*absolutely morally wrong; not at all morally acceptable; not at all morally required*), to 4 (*neither right nor wrong; somewhat acceptable; somewhat required*), to 7 (*absolutely morally right; completely morally acceptable; absolutely morally required*).

Table 4.4. Descriptive Statistics of Tipping Points (Study 4.3)

Species Killed	Species Saved	Never right to kill 10, no matter how many saved	Minimum	To save 100, never right to kill	Maximum number
			number saved for it to be right to kill 10 (Median)		killed for it to be right to save 100 (Median)
Human	Human	46.5%	30	7.8%	5
	Pig	83.8%	100	6.4%	0
Pig	Human	14.8%	1	49.6%	70
	Pig	18.3%	15	14.9%	50

Discussion

The results of this study confirm my hypothesis. People generally consider it to be permissible to harm animals irrespective of who will be saved—with only minor differences in moral judgments depending on whether humans or animals would be saved. At the same time, however, people consider it morally wrong to harm humans for animals, but less wrong if humans

can be saved. Harming humans to save more humans—as in the traditional Footbridge case—remains controversial among participants. Therefore, as I predicted, the difference in how permissible people consider it to harm humans if either humans or animals can be saved is much greater than the difference in how moral people consider it to harm animals if either humans or animals can be saved.

My explanation is that deontological constraints against harming animals are more easily overruled by utilitarian cost-benefit analysis than deontological constraints against harming humans. Further, in the case where humans have to be harmed to save many animals even cost-benefit analysis advises against the harmful action.

I found that there was an interaction between the species-membership of the individuals that have to be harmed and participants' individual level of speciesism on moral judgments, but no similar interaction between the species-membership of the individuals that can be saved and speciesism. This suggests that speciesism could be an underlying driver of the weaker deontological constraints people have for animals than humans. This is in line with Multi-level Weighted Deontology since speciesism measures the degree to which people perceive animals to have lower moral status than humans.

Study 4.4. Failed Attempt of Speciesism Reduction

The fact that people consider it more permissible to harm individuals of certain species for the greater good but not those of others might suggest that speciesism could be an underlying psychological driver of the strength of deontological constraints against instrumental harm. This is supported by the fact that individual differences in speciesism correlate with people's moral judgments of harming animals (but not human) for the greater good. To test this hypothesis more directly, I conducted a study in which I aimed to experimentally test the causal role of speciesism on instrumental harm.

Methods

Participants. 600 US American participants took part in the study online via MTurk and received \$1 payment for their participation. 162 were excluded for failing at least one of two attention checks, leaving a final sample of 438 people (205 female; $M_{age} = 36.50$, $SD = 11.65$). The reason that many participants were excluded because I included an attention check question that could only be answered if one watched the video fully.

Design, Materials and Procedure. The study had a 2 speciesism-reduction (yes vs. control) x species (pigs vs. humans) between-subjects design. In the speciesism-reduction condition participants read an article that described the complex inner and social lives of animals and included photos of different animals. It concluded with explicit moral arguments against speciesism. Next, participants had to do a short writing task that involved noting down the best reasons that animals deserve protection from cruelty. Finally, participants watched a 90-seconds video that showed animals. In the control condition, participants completed analogous tasks that focused on mobile phones and were unrelated to animals.

Results

A two-way ANOVA revealed that participants considered it equally permissible to harm pigs or humans to save more pigs or humans respectively when in the speciesism-reduction condition (humans: $M = 2.91$, $SD = 1.96$; pigs: $M = 4.86$, $SD = 1.68$) compared to the control condition (humans: $M = 2.76$, $SD = 1.70$; pigs: $M = 4.88$, $SD = 1.70$), $F(1, 434) = 0.15$, $p = .70$, $\eta^2 < .001$ (Figure 4.5). Similarly, there were no significant differences in the levels of speciesism between the speciesism-reduction ($M = 2.87$, $SD = 1.34$) and control conditions ($M = 3.10$, $SD = 1.23$), $t(427) = 1.78$, $p = .07$, $d = .17$.

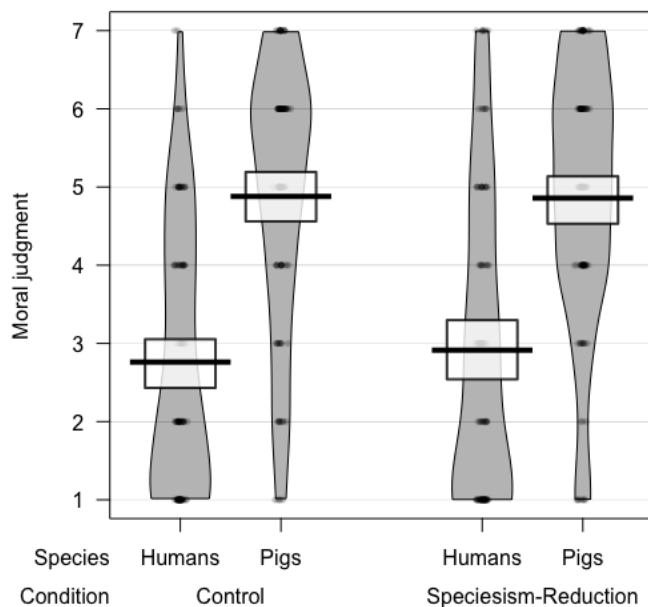


Figure 4.5. Moral judgments about harming few humans or pigs to save more, ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*). There were no differences between the speciesism-reduction and control conditions.

Discussion

In this study I aimed to experimentally manipulate speciesism to observe a change in the extent to which people hold deontological constraints for animals. However, I failed to manipulate speciesism as well as the strength of deontological constraints people hold for animals. I conclude that a person's degree of speciesism cannot easily be influenced. This adds to my previous finding of Chapter 2 that speciesism is a stable attitude that stays constant over a prolonged period of time. Therefore, I have to further rely on correlational approaches to gain some insight into the role speciesism plays in modulating the strength of deontological constraints.

Study 4.5. Between Objects and Humans

So far I have found that people consider it more permissible to harm animals than humans for the greater good. Does this mean that people see animals just as objects that can be sacrificed for any reason? In this study I aimed to compare the extent to which people hold deontological constraints for humans, animals, and objects. In particular, my research question is whether moral judgments of harming a few animals for many is closer to moral judgments of harming a few humans to save many or moral judgments of harming (or destroying) a few objects to save many.

Nozick (1974) hypothesised that people view animals as belonging into a moral category between humans and objects. On the one hand, he assumed, in some ways we treat animals like objects. We consider it permissible to own animals or to harm them for the greater good—something most people now consider to be wrong if it involved humans. On the other hand, since animals can suffer, there are many things we are not allowed to do with animals that we are allowed to do with objects. For example, while it may be permissible to harm animals for the greater good, it is generally not permissible to unnecessarily harm animals for weaker reasons such as for pure personal pleasure. With objects, however, people appear to mostly consider it permissible to do with them whatever one wants as long as nobody gets hurt or nobody's property rights are violated. A study by Mallon & Nichols (2006), however, found that people apply very weak deontological constraints to objects. It is not obvious, therefore, how deontological constraints for objects and animals compare.

My hypothesis, which I pre-registered at <https://osf.io/aeu3g/>, was that people's moral judgments of harming animals lies somewhere between their moral judgments of harming humans and objects. This is because, according to the Multi-level Weighted Deontology model, I hypothesise that people have extremely weak deontological constraints for objects, somewhat weak constraints for animals and strong constraints for humans.

The philosopher Martha Nussbaum (1995) has identified a number of properties of objectification such as ownership, instrumentality, denial of autonomy or fungibility (interchangeability). Based on this, I hypothesised these are applied to objects as well as, to some extent, to animals. People are likely to agree that objects are interchangeable, i.e. that one could easily replace one chair or toothbrush by another if they are similar and that they can be owned. The same does not apply to humans since every human is seen as being unique, whereas animals might be in a moral category in between humans and objects. Similarly, people are likely to agree that one object is less valuable than two, i.e. that one can say which of these two options is better. In contrast, from previous research we know that people are often unwilling to compare different options involving humans, and therefore they are unwilling to agree that one human is less valuable than two. Again, I hypothesise that animals are in an intermediate moral position. Overall, I expect that participants' moral judgments of harming a few to save many might track their views on interchangeability and comparability of these objects. This is because harming a few to save many only is a permissible option if one believes that the respective entities are in principle interchangeable and the two options comparable.

Method

Participants. I recruited 603 US American participants online via MTurk who received \$0.72 payment for their participation. Forty-six were excluded for failing at least one attention check, leaving a final sample of 557 people (299 female; $M_{age} = 38.39$, $SD = 11.77$). I expected a small to medium effect size. My analysis showed that 576 participants were required to detect an effect size of $f = 0.13$ with an alpha of 0.05, power of 0.80, and 3 groups (Cohen et al., 2014; Soper, 2017). I aimed to recruit 600 participants to account for any exclusions.

Design, Materials and Procedure. This study had a between-subjects design with three conditions: humans, animals, and objects. Each participant received two vignettes in random order. The reason I chose two was because I knew from Study 4.2 that there can be differences

between different types of animals. Similarly, I believed that there could be small differences between high and low value objects (cf. Mallon & Nichols, 2006). The two vignettes participants received involved adults and children in the human condition, dogs and pigs in the animals condition, and paintings and pens in the objects condition. I was primarily interested in the average effects, which is why I aimed to take the mean of the responses to the two vignettes in each condition as the main dependent variable. The vignettes were based on my vaccination vignettes from the previous studies but were slightly adapted to ensure plausibility and consistency across the three conditions. In all vignettes a few entities had to be harmed (e.g. killed or destroyed) to save many more of the same type from being either killed or destroyed.

After completing the main task, which involved indicating their moral judgments about harming a few to prevent similar harm to many in two separate vignettes, participants were presented with short additional questions. First, participants were asked whether they consider it permissible to kill a pig or a human/destroy a chair if one has a strong personal preference to do so. Second, they were asked whether they consider it permissible to own a chair, pig or human. Third, they were asked whether they think a chair, pig or human can easily be replaced by another one (interchangeability). Fourth, they were asked whether they believe one chair, pig, or human is less valuable than two (comparability). Next, participants were explicitly asked to indicate on a slider what moral value they attribute to animals relative to objects (-50) and humans (+50). Finally, as in the previous studies, participants then completed the Speciesism Scale, the Oxford Utilitarianism Scale, and responded to demographic items.

Results

Moral judgment about the two vignettes within each condition correlated strongly with each other (humans: $r = 0.95, p < .001$; animals: $r = .87, p < .001$; objects: $r = .74, p < .001$). Hence, the responses of the two vignettes in each condition were averaged to form a single score. Further

analysis revealed that if the two vignettes were analysed separately, the overall pattern of the results were the same. Therefore, the following analyses are based on the averaged scores only.

A one-way ANOVA revealed that there were significant differences across the three conditions in participants' moral judgments $F(2, 554) = 105.67, p < .001, \eta^2 = .28$ (Figure 4.6; Table 4.5). People considered it permissible to destroy a few objects in order to prevent more objects from being destroyed ($M = 5.83, SD = 1.28$). And as in my previous studies, most people considered it permissible, although less than for objects, to harm a few animals to save many animals ($M = 4.70, SD = 1.66$) but morally wrong to harm a few humans to save many humans ($M = 3.28, SD = 2.00$). Tukey's HSD post-hoc analyses revealed that responses in all three conditions significantly differed from each other ($p < .001$ for both comparisons with animals).

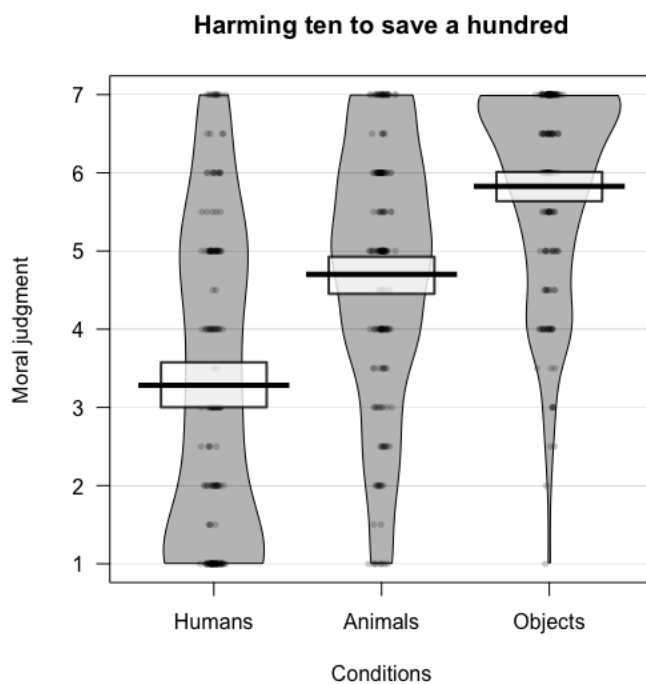


Figure 4.6. Moral judgments about harming few humans, animals, or objects to save more. Moral judgments of harming (or destroying) 10 humans, animals or objects to save 100 humans, animals or objects respectively, ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*).

Table 4.5. Moral Judgment and Additional Measures (Study 4.5)

	Humans	Animals	Objects
Harming a few to save many	3.28 (2.00)	4.70 (1.66)	5.83 (1.28)
Harming for personal preference	1.56 (1.23)	3.16 (1.72)	4.95 (1.53)
Owning	1.43 (1.07)	5.76 (1.24)	6.47 (0.9)
Replacing with another (interchangeability)	1.75 (1.19)	4.46 (1.7)	5.98 (1.03)
One less valuable than two (comparability)	2.35 (1.66)	4.48 (1.83)	5.36 (1.60)

Note. M (SD) ratings on a scale from 1 (*strongly disagree*) to 7 (*strongly agree*) or 1 (*morally wrong*) to 7 (*morally right*) respectively

Next, I looked at the other dependent measures. For all four questions I found the same pattern, with animals being in a category between humans and objects (Table 4.5). First, participants thought it was permissible to destroy a chair if one has a strong preference to do so, but it was somewhat wrong to kill a pig if one has a strong preference to do so, and very wrong to kill a human if one has a strong preference to do so. Second, participants thought it was permissible to own a chair, as well as a pig—even though somewhat less—but they thought it was wrong to own a human. Third, they thought that a chair could easily be replaced by another one, whereas a human could not be replaced by another human. For pigs, they thought it was somewhat possible to replace a pig with another one. Fourth, participants agreed that one chair was less valuable than two, slightly agreed that one pig was less valuable than two, but strongly disagreed that one human was less valuable than two. A regression model predicting moral judgments of harming a few to save many including condition, interchangeability ($b = .31$) and comparability ($b = .22$) as predictors revealed that both interchangeability and comparability were significant predictors and condition ceased to be significant after including the two, $F(4, 552) = 93.61, p < .001, R_2 = .40$. When the variable ‘owning’ was included in the model it was not a significant predictor.

When asked explicitly what the moral value of animals relative to humans and objects were, participants stated that they consider animals to be closer to humans with a mean of 23.68 ($SD = 23.54$) on a scale from -50 (objects) to +50 (humans), which was significantly above the mid-point 0, $t(556) = 23.74, p < .001, d = 1.01$. As in previous studies, the more speciesist participants were, the more permissible they considered harming animals ($r = .36, p < .001$), whereas the correlation between speciesism and harming objects ($r = .15, p = .05$) was not significant. The correlation between speciesism and harming humans was significant ($r = .14, p = .05$). The correlation coefficients between the animals and humans conditions differ significantly, $z = 2.28, p = .02$, as well as between the animals and objects conditions, $z = 2.24, p = .02$. As in previous studies, the more strongly participants endorsed instrumental harm (OUS-IH), the more permissible they considered harming any type to save more, whereas for endorsement of impartial beneficence (OUS-IB) there were no clear patterns of correlations. Younger people ($r = -0.30, p < .001$) and men ($r = -0.38, p < .001$) considered it more permissible to harm individuals to save more such individuals.

Discussion

This study showed that animals are perceived to lie in a moral category between objects and humans. People have very strong deontological constraints for humans, relatively weak ones for animals, and very weak ones, at best, for objects. These findings support the Multi-level Weighted Deontology model. They suggest that there is a hierarchy of moral levels, and that deontological constraints do not fully disappear as we move down the levels but just get weaker.

Similar patterns were found on views about harming for idiosyncratic reasons, ownership, interchangeability, and comparability. People consider it wrong for humans to be owned, but permissible for animals and even more for so objects to be owned. They considered it very wrong to harm humans to fulfil a personal preference, fairly wrong to harm animals but permissible to harm objects to fulfil a personal preference. Furthermore, I found deontological constraints against

harming for the greater good may be related with—or even underpinned by—views on interchangeability and comparability of the respective entities. People disagree that humans can easily be exchanged by other humans or that one human life is less valuable than two. At the same time, they have completely different views when it comes to objects: they believe that objects can easily be replaced by other objects and that one object is clearly less valuable than two objects. Again, animals are seen in an intermediate position with regard to these two aspects. My results suggest that only if one believes that an entity can be compared with and replaced by other similar entities, is it considered sensible to harm one such entity to save many.

Study 4.6. Mental Capacities

In all previous studies I found that people have weaker deontological constraints for animals than for humans. I have also assumed that according to Multi-level Weighted Deontology people's deontological constraints get reduced, the less they morally value the individuals at stake. However, this leaves open the question of what grounds people have for the attribution of greater or lesser moral status to these individuals. One factor which I have identified is speciesism, namely that people attribute lower moral status to animals solely due to their species-membership. This hypothesis is supported by my consistent finding across my studies that individual differences in speciesism correlate with weaker deontological constraints for animals but not humans.

In this study, I want to examine the possibility that there are additional factors that are associated with (but distinct from) species-membership that underpin deontological constraints, namely the perceived levels of mental capacities. In Chapter 3 I showed that perceived mental capacities of animals can to some extent contribute to moral anthropocentrism. In this study, I will in particular investigate whether perceived cognitive capacities and sentience of individuals influence the extent to which people grant them deontological protections against harm.

Advanced cognitive capacities enable an individual to have language, complex thoughts, deliberate judgments and autonomy—abilities that are often considered a requirement to be considered a rational person and moral agent. Having sentience enables an individual to experience suffering, happiness, and other emotions. Both factors are central properties of mind perception, which some have claimed to be at the heart of moral status (Gray et al., 2012; Piazza, Landy & Goodwin, 2014). Humans and animals are generally both considered sentient. But while animals have only basic cognitive capacities, humans generally have more advanced cognitive capacities. Plants or inanimate objects are lacking both sentience and any form of cognitive capacity. A sophisticated robot might have advanced cognitive capacities but is not sentient. Can these two properties (or the combination of them) explain the difference in strength of deontological constraints for humans and animals above and beyond pure species-membership?

To test this, I presented participants with scenarios involving a hypothetical extraterrestrial species who either had advanced or basic cognitive capacities and who either had a low or a high degree of sentience. In a similar study Sytsma and Machery (2012) asked participants whether it would be wrong for humans to experiment on or kill hypothetical individuals to benefit humans. They found that participants considered it more wrong if the individuals had advanced compared to basic cognitive capacities. However, they found no effect of sentience nor an interaction effect. Their study, however, did not ask whether it would be permissible to harm such individuals to save more individuals of the same type. I, in contrast, was primarily interested in whether participants consider it wrong to harm a few hypothetical individuals to save many more individuals of the same type. My hypothesis, which I pre-registered at <https://osf.io/24vwr/>, was that participants consider harming individuals with advanced cognitive capacities (just like humans) to be wrong but individuals with basic cognitive capacities (just like animals) to be generally permissible. Similarly, I hypothesised that participants consider harming sentient individuals to be more immoral than non-sentient individuals. I did not have a clear hypothesis whether there would

be an interaction effect between cognitive capacities and sentience. Furthermore, I asked whether participants considered it wrong to harm a few individuals for weak reasons, such as for small personal benefits like for tastier (but equally healthy) food purposes. By doing so, I aimed to test the hypothesis that people consider it permissible to harm non-sentient individuals with basic cognitive capacities even for small benefits.

Method

Participants. I recruited 812 US participants online via MTurk, and they received \$0.5 payment for their participation. One hundred and eighty-seven were excluded for failing at least one attention check or manipulation check, leaving a final sample of 625 people (307 female; $M_{age} = 35.72$, $SD = 11.30$). I expected a small effect size. My analysis showed that 787 participants were required to detect an effect size of $f = 0.1$ with an alpha of 0.05, and power of 0.80 (Cohen et al., 2014; Soper, 2017). I aimed to recruit 800 participants to account for any exclusions.

Materials and Procedure. The study had a 2 cognitive capacities (basic vs. advanced) x 2 sentience (low vs. high) study design with both factors being manipulated between subjects. The reason I chose a hypothetical species was to avoid participants being influenced by stereotypes they might hold about, and beyond, the cognitive capacities and sentience of familiar species. This study design and stimulus was loosely based on the study by Sytsma and Machery (2012), as well as aspects of Starmans and Friedman (2016). Participants read a fictional paragraph about a planet on a nearby solar-system on which extraterrestrial individuals called ‘Atlans’ are the dominant form of life. Participants then read one of four descriptions of the Atlans depending on condition. The Atlans were described as either having a high degree of sentience (“feel pleasure and pain as well as other emotions such as fear, anger and sadness”) or as having a low degree of sentience (“don’t feel pleasure or pain or any other emotions such as fear, anger, or sadness”). Further, they were described as either having advanced cognitive capacities (“have opinions, beliefs, desires, language,

and culture. [...] can make decisions for themselves. [...] make plans for their long-term future”) or as having basic cognitive capacities (“don’t have opinions, beliefs, desires, language, or culture, [...] respond to their environment purely based on impulse and instinct [...] don’t make plans for their long-term”).

Following these descriptions, participants were given the Vaccine Death scenario from Study 4.2b, according to which ten healthy Atlans needed to be killed to develop a vaccine to save a hundred other Atlans from dying of a disease. Next, participants were given two additional scenarios for exploratory purposes. In one ten Atlans had to be killed to save a hundred humans, and in the other Atlans had to be killed to produce a tasty food alternative.

Following these three scenarios, participants responded to manipulation check questions, which tested whether they correctly understood and remembered what cognitive capacity level or sentience level the extraterrestrials had. Next, participants completed a range of other measures. They answered three questions about the Atlans’ ability to consent. As in the previous studies, participants then completed the Speciesism Scale, the Oxford Utilitarianism Scale, and responded to demographic items.

Results

A two-way ANOVA revealed that participants considered it more permissible to harm ten individuals to save a hundred individuals when these individuals had basic cognitive capacities (sentient: $M = 4.33$, $SD = 1.74$; not sentient: $M = 4.44$, $SD = 1.71$) compared to if they had advanced cognitive capacities (sentient: $M = 3.90$, $SD = 1.81$; not sentient: $M = 3.95$, $SD = 1.83$), $F(1, 621) = 10.56$, $p = .001$, $\eta_p^2 = .02$ (Figure 4.7; Table 4.6). However, the level of sentience made no difference on moral judgements, $F(1, 621) = 0.34$, $p = 0.56$, $\eta_p^2 = .001$, nor was there an interaction between cognitive capacities and sentience $F(1, 621) = 0.06$, $p = .80$, $\eta_p^2 < .001$.

The same pattern was found for the the next question. Participants considered it more permissible to harm ten individuals to save a hundred humans when the individuals had basic

cognitive capacities (sentient: $M = 3.81$, $SD = 1.93$; not sentient: $M = 4.15$, $SD = 1.89$) compared to if they had advanced cognitive capacities (sentient: $M = 3.10$, $SD = 1.98$; not sentient: $M = 3.03$, $SD = 1.95$), $F(1, 621) = 34.68$, $p < .001$, $\eta_p^2 = .053$. Again, the level of sentience made no difference on moral judgements, $F(1, 621) = 0.81$, $p = .37$, $\eta_p^2 = .001$, nor was there an interaction between cognitive capacities and sentience, $F(1, 621) = 1.81$, $p = .18$, $\eta_p^2 = .003$.

Next, I looked at moral judgments about killing individuals for food. Participants generally considered it very immoral to harm individuals for small personal benefits such as tasty food. Participants considered this more immoral for individuals with advanced cognitive capacities (sentient: $M = 2.27$, $SD = 1.73$; not sentient: $M = 2.38$, $SD = 1.78$) than individuals with basic cognitive capacities (sentient: $M = 1.43$, $SD = 1.12$; not sentient: $M = 1.44$, $SD = 1.05$), $F(1, 621) = 57.97$, $p < .001$, $\eta_p^2 = .09$, whereas sentience, again, made no difference, $F(1, 621) = 0.25$, $p = .61$, $\eta_p^2 < .001$, neither was there an interaction, $F(1, 621) = 0.18$, $p = .67$, $\eta_p^2 < .001$.

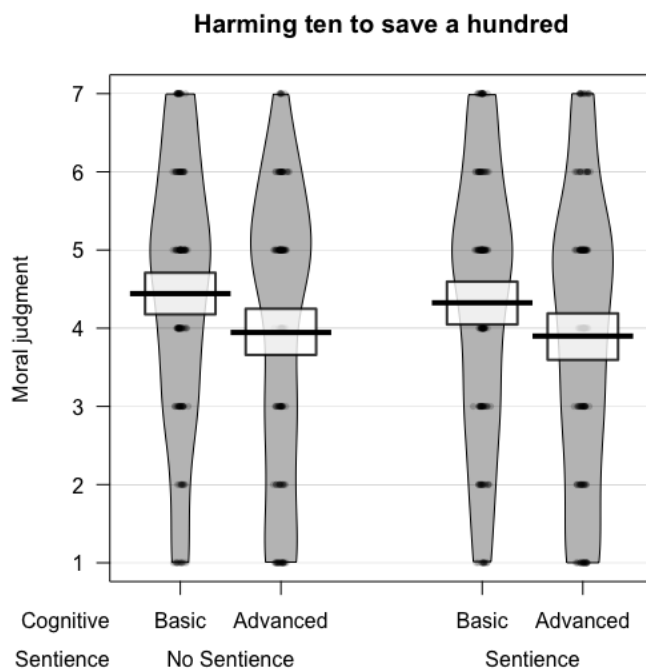


Figure 4.7. Moral judgments of harming 10 individuals of a hypothetical species to saving 100 individuals of the

same species, ranging from 1 (*absolutely morally wrong*), to 4 (*neither right nor wrong*), to 7 (*absolutely morally right*).

The more speciesist participants were, the more permissible they considered harming individuals, irrespective of the reason for why the individuals were being harmed (Table 4.6). Similarly, the higher participants scored on instrumental harm (OUS-IH), the more permissible they considered harming individuals for the greater good. Furthermore, I found that the more participants believed the individuals were able to give consent, the more immoral they considered it to harm them. Younger participants considered it more permissible to harm individuals to save more such individuals, $r = .20, p < .001$, and men considered it more permissible as well, $r = -.16, p < .001$. There were no other noteworthy correlations with demographics.

Table 4.6. Correlations with Moral Judgment of Harming Hypothetical Individuals to Save Many Across Conditions (Study 4.6)

<i>Cognitive Capacities</i>	<i>Sentience</i>	Speciesism	Consent	OUS-IH
<i>Advanced</i>	<i>High</i>	.21**	.12	.52**
	<i>Low</i>	.16*	-.11	.56**
<i>Basic</i>	<i>High</i>	.19*	.21**	.53**
	<i>Low</i>	.14	-.17*	.38**

Note. * $p < .05$. ** $p < .01$. Consent represents to what extent participants thought the harmed aliens could consent. OUS-IH represents to what extent participants endorsed instrumental harm for the greater good.

Discussion

This study showed that an individual's cognitive capacity level is a significant, but small, determinant of the strength of the deontological constraint people grant that individual. People consider it less permissible to harm individuals with advanced cognitive capacities than individuals

with basic cognitive capacities. An individual's level of sentience, in contrast, was not a significant determinant.

Since animals generally have basic cognitive capacities, but humans generally advanced ones, it is possible that part of the explanation for why people hold deontological constraints for humans but not animals is due to the different cognitive capacity levels. However, it should be noted that the effect size of cognitive capacities on moral judgments was very small ($\eta^2 = .02$), and much smaller than the effect sizes I observed in my previous studies, which explicitly contrasted animals and humans ($\eta^2 = .14$ in Study 4.2). This suggests that while cognitive capacities explain part of the effect, species-membership itself likely explains the largest part of it. This is in line with my previous research in which I found that speciesism (i.e. discrimination based purely on species-membership) explained stronger helping behaviour towards humans compared to animals above and beyond perceptions of level of cognitive capacities (or intelligence) and sentience (Chapter 2).

The fact that the level of sentience made no impact on participants' moral judgments was somewhat surprising considering that reduced perceptions of sentience are associated with perceptions of lower moral status of animals (e.g. Chapter 2; Gray et al., 2012). It is also inconsistent with the finding that people are less likely to harm robots to save many humans if the robots are described as having high degree of affective experience than low degree (Nijssen et al., 2019).

One possibility is that my manipulations were suboptimal. For example, my sentience manipulation only focused on the ability to experience hedonic states and emotions, but not on an individual's level of consciousness. Furthermore, it is possible that my cognitive capacities manipulation was slightly confounded with elements of sentience since it described 'desires', which are mental phenomena that could be seen as aspects of sentience.

Overall though, I still believe that these results support the view that an individual's level of sentience does not strongly influence the deontological protection people grant that being as

long as it has at least some cognitive capacities. In fact, this finding is in line with Sytsma and Machery (2012) who also found no difference between people's judgments of harming sentient vs. non-sentient individuals to benefit humans. My findings also broadly fit with the findings of Cikara et al. (2010) who found that a person's level of competence (arguably related with cognitive capacities) but not of warmth (arguably related with sentience) has an influence on the strength of deontological constraints. Finally, note that I excluded participants who did not correctly indicate the Atlans' level of sentience in a manipulation check.

The fact that participants considered it immoral to harm all of the different types of hypothetical individuals for tastier food suggests that they did not conceptualize the extraterrestrials as (food) animals or plants in any of the conditions but rather as human-like beings. It is possible that people consider harming individuals for food purposes only morally permissible for specific animals such as pigs, chickens and cows and not for all non-human beings (cf. meat paradox; Loughnan, Haslam, & Bastian, 2010)

Study 5.7. Time Investment and Donations

The previous studies all focused on judgments in hypothetical scenarios and were presented to online samples on MTurk. In this last study I aimed to replicate the effect with an offline student sample with two goals: first, to show that the effect shows in the real-world context of medical experimentation; second, to show that the hypothetical judgments I studied so far translate into behaviour. In the real world, harming animals to benefit others is common in the context of medical experimentation. The general societal consensus is that medical experimentation on animals, even if harmful, is justified. Whereas on humans it is not—in particular if no consent by the humans is given.

In this study, I asked students on the University of Oxford campus about their opinion on the ethical permissibility of medical experiments on pigs or human infants. In addition to

measuring their opinions in the form of judgments, I also aimed to measure behaviour. For that, I gave them the possibility to invest their personal time to help improve the campaign of an activist organization that was fighting against the proposed medical research program on pigs or infants. Furthermore, participants had the option of donating personal money to support the activist organization if they wished to do so. My hypothesis was that I would replicate the effect found in the previous studies for judgment as well as for behavioural measures.

Method

Participants. I recruited 208 students on the campus of the University of Oxford who received £3 payment (in line with UK minimum wage) for their participation. Eight were excluded because they did not complete an attention check or did not finish the study, leaving a final sample of 200 people (118 female; $M_{age} = 24.31$, $SD = 7.68$). Power analysis showed that 197 participants were required to detect an effect size of $f = 0.2$, which I determined in an online pilot study, with an alpha of 0.05, and power of 0.80 (Cohen et al., 2014; Soper, 2017). I aimed to recruit 200 participants to account for any exclusions.

Materials and Procedure. The study had two between-subjects conditions. Participants first read a text about a planned medical experiment at the university's medical science lab. The research was described to involve experimentation on either 50 young pigs or human infants depending on the condition in order to develop a medicine that would help thousands of sick pigs or human infants. The text stated that while the research could be painful for the test subjects, it would not have any long-term negative side effects. Next, participants were informed that due to ethical concerns of this planned medical research program the advocacy group *Unethical Research Watch* plans to fight against the implementation of the research. In order to do so, the advocacy group asked us to conduct a survey to assess the opinions of the general public, whose results will inform their campaign. As part of this alleged survey, participants were asked to indicate whether they considered the planned research program to be ethically justified or not, and whether they

would like it to stop. On the next page, participants were told that if they shared the advocacy group's view and wanted to support their campaign, they could do so by completing a short writing task. In particular, participants were told that the advocacy group was looking for the best arguments to support their view, which they could write down on the paper if they wished to support the campaign, and that I would forward their responses to the advocacy group in order for them to improve their campaign. Finally, participants were also given the opportunity to donate any amount between 0 and 300 pennies of their £3 payment to support the advocacy group.

Participants also completed the Speciesism and Oxford Utilitarianism Scales, as well as demographic questions. At the end of the study, participants were debriefed about the fact that both the planned medical research program as well as the advocacy group were made up. Participants kept the money they intended to donate.

Results

Participants considered the proposed medical research program significantly more ethically justified if it involved pigs ($M = 5.39$; $SD = 1.34$) compared to human infants ($M = 4.34$; $SD = 1.77$; *Figure 4.8*), $t(182.58) = -4.71, p < .001, d = 0.67$. Similarly, they were significantly more likely to say that they wanted the program to stop if it involved infants ($M = 3.49$; $SD = 1.88$) compared to pigs ($M = 2.69$; $SD = 1.50$), $t(186.76) = 3.35, p < .001, d = 0.48$. Next, I looked at the behavioural measures. Significantly more participants were willing to invest time to help to support the campaign of the activist organization fighting against the medical research program if the research involved human infants (38%) than pigs (19%), $\chi^2(1) = 7.71, p = .005$. Finally, participants also donated significantly more money (pennies) to support the activist group if the research involved human infants ($M = 51.52$; $SD = 107.96$) than pigs ($M = 21.15$; $SD = 70.27$), $t(168.85) = 2.34, p = .02, d = 0.33$. Since donations were not normally distributed, I conducted a Wilcoxon rank test, which supported the finding, $W = 5348.5, p = .03$.

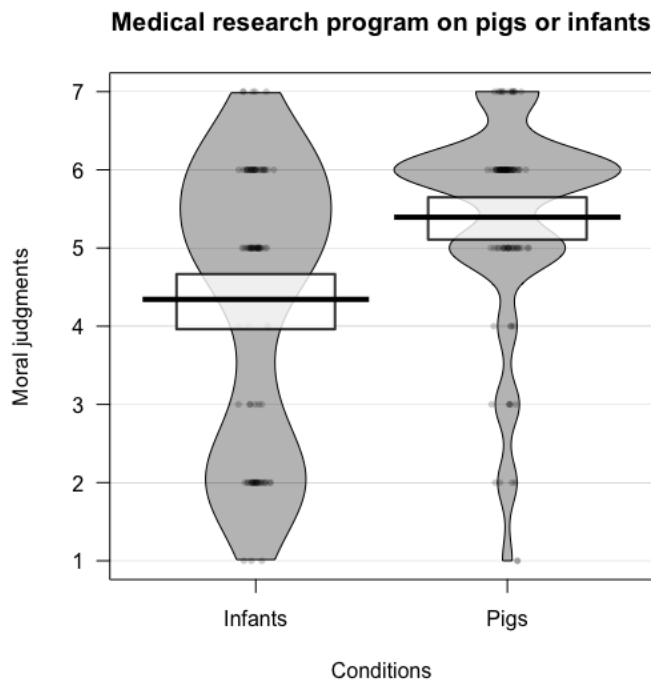


Figure 4.8. Judgment about the ethical justifiability of the medical research program that harms pigs to save many more pigs or human infants to save many more human infants, ranging from 1 (*completely unjustified*), to 4 (*neither justified nor unjustified*), to 7 (*completely justified*).

As in the previous studies, the more speciesist participants were, the more ethical they considered the medical research program on pigs, $r = .45, p < .001$, but not the medical research program on infants, $r = .16, p = .11$. The two correlation coefficients differed significantly, $z = 2.25, p = .02$. Endorsement of instrumental harm (OUS-IH) correlated significantly with ethical judgments in both the pigs, $r = .26, p = .01$, as well as infants, $r = .51, p < .001$ conditions. Endorsement of impartial beneficence (OUS-IB) did not correlate with ethical judgments. Younger participants thought the research was more ethical than did older participants ($r = -.17, p = .02$). There was no significant correlation of age with moral judgement.

Discussion

This study replicated the effect I found in the previous studies in an offline student sample, namely that people have weaker deontological constraints for animals than for humans. The study demonstrated that the effect shows in the real-world context of medical experimentation and, crucially, that it translates into behaviour—or at least decision-making with assumed real world consequences. People were more willing to invest both their time and personal money to support an activist campaign against medical research on human infants than on pigs. Nevertheless, there was some degree of opposition to research on animals even when this would help thousands of other animals. This is in line with the Multi-level Weighted Deontology model that suggests that deontological constraints get weaker, but do not disappear, the lower the perceived moral status of the respective individual. As in the previous studies, reduced levels of speciesism were associated with weaker deontological constraints for animals. This suggests that speciesism does not just capture the tendency to prioritise individuals of certain species over others, but also the tendency to hold weaker deontological constraints for individuals of certain species, the lower they are in the moral hierarchy.

It is worth noting that in this study the cognitive capacity level in humans and animals was being held constant since both human infants as well as young pigs only have basic cognitive capacities. Yet, I still observed a difference in strength of deontological constraints. One explanation could be that people perceive human infants to have personhood despite not having advanced cognitive capacities themselves because they have the potential to gain advanced cognitive capacities or because they are members of a species whose typical members have advanced cognitive capacities (e.g. Cohen, 1986; Scanlon, 1998; Finnis, 1995; Kagan, 2016). An alternative hypothesis is that species-membership itself, i.e. speciesism, plays a key role in explaining the effect.

General Discussion

My studies show that people hold much weaker deontological constraints against harming animals than against harming humans. In this chapter's studies I found that people consider it more permissible to harm animals to save many animals than to harm humans to save many humans. The effect is strong and showed in different types of dilemmas: in personal (Footbridge case) and impersonal (Switch case) sacrificial dilemmas (Study 4.1), and in both abstract (Study 4.3) and concrete dilemmas (e.g. vaccine testing; Study 4.2b, 4.5). The effect also showed behaviourally in the real-world context of opposition to medical experimentation (Study 4.7). The more speciesist people are, the weaker their deontological constraints against harming animals. And the less they value a certain type of animal, the weaker the deontological constraints against harming that animal (Study 4.2b). Constraints for chimpanzees and dogs were stronger than constraints for pigs. But by far the strongest difference was that between deontological constraints for humans and for animals in general. Deontological constraints against harming animals were weak but never disappeared completely. And they were stronger than constraints against damaging objects (Study 4.5). I found that primarily species-membership was driving the strength of the deontological constraints. Cognitive capacities played a weak additional role, whereas sentience played no role (Study 5). These findings are in line with my findings of Chapter 2 that showed that species-membership has a stronger influence on moral status than cognitive capacities and sentience.

These findings fit my hypothesis of Multi-level Weighted Deontology and rule out the competing accounts of people's ethical views about harming humans and animals that I outlined in the introduction of this chapter. Utilitarianism and Cross Species Deontology can be ruled out because people generally value humans more than animals. Kant's view can be ruled out because people do value animals at least to some extent (although this could also be due to instrumental reasons). Multi-level Uniform Deontology does not take into account the fact that deontological constraints against harming animals are weaker than those against harming humans even if many

of the same respective beings can be saved. Nozick’s “utilitarianism for animals, Kantianism for people”, while being a good first approximation, does not account for the fact that people *do* have weak deontological constraints for animals and therefore are not completely “utilitarian” for animals. Furthermore, Nozick’s slogan does not account for the fact that deontological constraints for different types of animals differ depending on how much people value them, e.g. constraints against harming dogs are stronger than those against harming pigs. Multi-level Weighted Deontology, in contrast, appropriately captures the fact that deontological constraints get weaker—without completely disappearing—the less people value the beings at stake. Of course, there were individual differences in participants’ responses, which means that for a small minority of people other models are more accurate descriptions. However, the majority of participants’ responses can best be captured by Multi-level Weighted Deontology.

Moral Judgments in Cross Species Sacrifice Cases

The cross-species sacrifice cases of Study 4.3 are revealing for our understanding of the psychological mechanisms at play. I found that people considered it only slightly more permissible to harm pigs to save many humans than to save many pigs. In contrast, people were much more sensitive to the type of individuals that could be saved if humans had to be harmed. They were strongly opposed to harming humans in order to save many pigs, whereas harming humans to save many humans was considered to be less wrong on average. How can this asymmetry be explained?

One way to explain this is as follows: The lower the moral status of the individual that needs to be harmed, the weaker the deontological constraint against harming that individual. The stronger the deontological constraints, the more difficult it is to overrule them. People have a rough threshold—that varies between individuals and contexts—for when to overrule a given deontological constraint. Whether this threshold is met depends on the utilitarian cost-benefit analysis, which people will engage in if they are capable of doing so. In cases where animals have to be harmed to save either many animals or many humans, the threshold is clearly met in both

cases—though since people value humans much more than animals, it is more easily met if humans instead of animals can be saved. In contrast, however, in cases where humans have to be harmed to save either many humans or many animals, the threshold is not clearly met, and the difference in how easily it is met is far greater than in the previous cases. The threshold for harming humans to save many humans may be met to some extent in some people. But since people value humans much more than animals, the threshold for harming humans to save animals will be met, if at all, only if a *vast* number of animals will be saved. As a consequence, the extent to which deontological constraints against harming are overruled differs across these four cases and results in the observed asymmetry.

This explanation assumes that people assign a certain value to the life of a single animal and another value to the life of a single human—all using the same moral value “currency”. These values are then aggregated each based on the number of individuals on each side and contrasted against each other. The contrasted aggregated values are then matched against a common threshold value that is measured in this shared currency. An advantage of this explanation is that it could also deal with mixed sacrifice cases, such as a case where 55 pigs have to be harmed to save 50 pigs plus 2 humans, which is a type of case I did not test but future research could look into.

A reason to be sceptical of the described view could be that it appears psychologically implausible that people engage in such complex calculations and rely on a common moral value currency. Instead, it may be more plausible that a much rougher estimate, that is itself partly emotional, is at work. One possibility is that the threshold for when deontological constraints should be overruled is different depending on what type of individuals can be saved. The threshold is lower, the lower the perceived moral status of the individuals that can be saved. This view, however, is harder to reconcile with the idea of a fixed aversive response against harming dictating the strength of the deontological constraint since it suggests that the constraint against harming

does not arise from an affective response to considerations of the harmful act in isolation, but is instead modulated by which individuals would be saved by that act, and how many. In other words, on this view the degree of aversion to harming and the calculation of consequences interact rather than being generated independently, as is often assumed by current models of moral judgement.

Implications for the Psychology of Moral Judgment

Current psychological research on deontological rules against harming others often models such rules on Kant's ethics, and therefore understands them as representing absolute prohibitions against certain actions. It is implausible, however, that lay people regard deontological rules as absolute in this way (Kahane, 2015). Lay people's intuitions accept a plurality of moral rules which can conflict in some cases. When this happens one rule will typically overrule another (e.g. we may break a minor promise if keeping it would be deeply unfair). Similarly, deontological constraints can be overruled when the consequences of following them are too severe—e.g. lying when this will prevent murder and, far more controversially, the use of 'enhanced interrogation methods' to prevent terrorist attacks. The present research provides a powerful demonstration of this central feature of common-sense morality. While the absolutist understanding of deontological rules suggest the all-or-nothing picture suggested by Nozick's proposal—absolute prohibitions protecting humans, cost-benefit analysis for animals—our findings strongly support a Multi-level Weighted Deontology account of judgments about harm towards humans and animals. I repeatedly found that deontological constraints are a matter of degree. And when and whether they will be overruled is a function both of the individuals to be harmed, the individuals to be saved, and the numbers of each.

My findings consequently also show that moral thinking about harm to humans and animals is not fundamentally different but rather varies in degree. People value animals much less than humans and accordingly grant them much weaker deontological constraints. Importantly, though it is not the case that people have no deontological constraints for animals whatsoever.

This difference in degree, however, can be very substantial, and most people consider it fairly wrong to harm humans while fairly permissible to harm animals for the greater good.

According to the dual-process model of moral judgment (Greene, 2014) an automatic and intuitive aversion to harm prevents more deliberate utilitarian cost-benefit analysis. Previous research has shown that in more personal moral dilemmas that involve direct actions (e.g. pushing a person off the bridge) the aversion to harm is more emotional and stronger compared to impersonal moral dilemmas that involve indirect actions (e.g. pushing a button; cf. Greene, 2014). Therefore, one possible conclusion is that harming individuals with lower perceived moral status, such as animals, is less personal and emotionally evocative, resulting in a weaker aversion to harming such individuals. Against this interpretation, however, speaks my finding of Study 1 that deontological constraints against harming animals and humans are reduced to a similar extent in both personal and impersonal moral dilemmas. This suggests that harming humans vs. animals is not analogous to personal vs. impersonal sacrificial dilemmas. If that were so, then we should expect an interaction effect in Study 4.1. Further research, however, is needed to investigate the role aversion to harm may play in explaining the differences I observed between the human and animal cases, and, more generally, to clarify the psychological mechanisms underlying thresholds for overruling deontological constraints.

To my knowledge these are the first studies that explicitly combine research on the psychology of speciesism (and moral attitudes towards animals more generally) with utilitarianism (as well as deontology). While deontological constraints, such as those against instrumental harm, are a key component of deontology, they only capture one side of (non-)utilitarianism. Another important and psychologically distinct aspect is impartial beneficence (Oxford Utilitarianism Scale; Kahane et al., 2018): a utilitarian endorses complete impartiality and is willing to help others irrespectively of who they are, i.e. also irrespectively of their species-membership (Singer, 1981). Conceptually, therefore, anti-speciesism can be seen as a subcomponent of impartial beneficence.

Indeed, my studies suggest that the two are also psychologically associated since speciesism and endorsement of impartial beneficence were strongly negatively correlated. At the same time, however, degree of endorsement of impartial beneficence was generally not correlated with endorsing stronger deontological constraints against harming animals in sacrificial dilemmas. This could reflect the difference between more abstract principles and judgment in concrete cases, but also the focus of the impartial beneficence subscale (Kahane et al. 2018) on *self*-sacrifice for the greater good, something not investigated here.

Animals Are Neither Objects nor Persons

My findings support Nozick's view (1974) that people place animals into a moral category between humans and objects. Animals receive some moral properties that resemble the ones that objects have: like objects, people consider it morally permissible to own animals, to treat them as a tool for another's purpose, to deny them of their autonomy, or to treat them as interchangeable with other animals—properties that the philosopher Martha Nussbaum (1995) has identified as aspects of objectification. However, in contrast to objects animals are seen as sentient beings that deserve *some* moral protection. People are, for example, against unnecessary deliberate animal cruelty (Vaughn, 2009). Similarly, I found that people consider it permissible to destroy an object but not to harm an animal, if the sole purpose is to fulfil a personal preference. (The case of meat consumption appears to be an exception and might be related to people's tendency to de-mentalize farmed animals; Bastian et al., 2012).

However, despite the fact that animals share moral features of objects, on other crucial dimensions animals are clearly not seen as objects. In contrast to objects, animals are seen as individuals that deserve *some* moral status. For example, I found that people consider it permissible

to destroy an object, but not to harm an animal, just to fulfil a strong personal preference⁶. People generally appear to be against unnecessary deliberate animal cruelty (Vaughn, 2009). In sum, perhaps one way of describing the moral status of animals is to say that animals are sentient objects.

The fact that deontological constraints for humans are much higher than those for animals is reflected by the fact that humans are seen as possessing inalienable rights, but such rights are rarely, if ever, ascribed to animals. Most people of Western societies are firm believers in basic human rights such as a right to life, bodily integrity and autonomy—rights that are inalienable and absolute. Even humans without advanced cognitive capacities such as human infants or mentally severely disabled humans are granted these same basic rights. Animals, in contrast, are not granted equivalent rights, not even the most basic ones. This shows, for example, in the context of medical experimentation or exploitation for consumption.

Limitations and Future Research

I have identified Multi-level Weighted Deontology as the model that best describes people's intuitions about when it is permissible to harm humans and animals. Future research could explore the model further and make it more precise. As discussed above, for example, the psychological mechanisms underlying thresholds for overruling deontological constraints are not clear yet. Future research could look at a wider range of cross-species sacrifice cases as well as at mixed cases in which a group of individuals consisting of different species would have to be harmed or saved.

I concluded that speciesism itself is a key factor that explains why deontological constraints for animals are much weaker than those for humans. While I was able to demonstrate this correlationally, I was not able to investigate the causal role of speciesism (as measured by the

⁶ The case of meat consumption appears to be an exception and might be related to people's tendency to deny the level of sentience of farmed animals (Bastian et al., 2012)

Speciesism Scale) directly. In Study 4.4 I failed to manipulate speciesism to observe its influence on deontological constraints. However, while I was not able to directly manipulate the individual level of speciesism, I did successfully manipulate species-membership (e.g. humans vs. pigs), which made a strong difference throughout all my studies. This, together with the finding that mental capacities only made a minor difference, suggests that speciesism is likely the primary driver of the strength of deontological constraints for animals.

It remains an open question which factors associated with the different species I tested are driving deontological constraints, or whether it is just species-membership per se. I tested cognitive capacities and sentience, but there could be several other factors that correlate with the species-membership such as whether society is used to killing certain animals or not, whether it is considered legal to kill certain animals or not, whether certain animals are frequently used in meat production or not, the animals' sociability, physical attractiveness, or other properties. Future research could disentangle these potential factors by testing a broader range of animals than the ones I have tested.

I found that people have weaker deontological constraints against harming animals than humans. Since animals are more distant to us than other humans, one may conclude that deontological constraints should be strongest for those who are closest to us. By contrast, Kurzban, DeScioli and Fein (2012) found that people are more willing to sacrifice their brother to save five brothers than they are to sacrifice a stranger to save five strangers. This suggests that deontological constraints are actually weaker for those closest to oneself. A possible way to explain this—but this is a matter for further investigation—is that holding deontological constraints against harming humans is especially important to ensure cooperation within human groups by signalling trustworthiness. Research has shown that people who do not hold such constraints (against harming humans) are considered untrustworthy (Everett, Pizarro, & Crockett, 2016; Everett, Caviola, Savulescu, & Faber, 2018). It is possible that there are two ways in which

deontological constraints could get reduced: First, in cases where the individuals in question are not considered to be part of the cooperative group, e.g. animals. Second, in cases where the agent's trustworthiness and cooperative intentions are not questioned in the same way, e.g. among family members. Note also that even though people care more about family members (i.e. exhibit partiality), they usually do not explicitly attribute higher moral status to family members than to strangers, whereas they do explicitly attribute lower moral status to animals than to humans.

A related open question is how Multi-Level Weighted Deontology relates to the different sub-factors of speciesism I uncovered in Chapter 3. Are Species-Relativist or Species-Absolutist factors responsible for the weaker deontological constraints against harming animals? One reason to assume that it is not Species-Relativism is the finding by Kurzban et al. (2012) mentioned above. If Species-Relativism is analogous to partiality towards one's family and people are more likely to harm one brother to save five brothers, one would assume that people should also be more likely to harm one human to save five humans compared to harming one animal to save five animals. Since that is not the case, this suggests that Species-Absolutist factors, such as Pro-Human Bias and Anti-Animal Bias, are responsible for the differential weighing of deontological constraints. In line with this, in Study 3.4 I found that the Anti-Animal Bias drives the strong tendency to harm an animal to save a human.

Conclusion

In sum, this chapter showed that deontological constraints against harming a few to save many more is not something that we have to the same extent for all types of individuals. Instead, the extent to which we hold such rules depends on the species-membership of the individuals at stake (i.e. speciesism). The less people value the individuals at stake, the weaker the deontological constraint against harming a few to save many. The vast majority of people consider harming animals for the greater good much more permissible than harming humans, as can be observed in

the context of medical experimentation. Crucially, though, deontological constraints against harm never completely disappear.

The findings of this chapter extend my findings of the previous chapters, by showing that speciesism is reflected in a crucial aspect of morality: deontological constraints against harm for the greater good.

5. Concluding Discussion

Humans regard themselves as having moral priority over all other animals, an attitude that manifests itself in numerous policies, practices and choices, ranging from consumption of animal meat and animal experimentations to zoos and animals performing in circuses. My aim in this thesis was to empirically investigate this attitude. I argued that a key driver of moral anthropocentrism is speciesism—the tendency to value animals less because of their mere species-membership.

Summary of Key Findings

Through the development of the Speciesism Scale in Chapter 2, I showed that speciesism is a stable psychological construct with high interpersonal differences and psychological associations with other forms of prejudice. Analogous to racism, sexism or homophobia, speciesism is linked to social dominance orientation—the view that group-based dominance and inequality can be justified. Beyond judgements, speciesism showed to be predictive of behaviour: the higher they scored on my speciesism scale, the more participants tended to help humans than animals (as well as certain animals such as dogs over others such as pigs). Importantly, my studies showed that speciesism can predict these behaviours above and beyond the perceived level of mental capacities people attribute to the respective beings. This supports the claim that species-membership itself is the key driver for differences in behaviour.

In Chapter 3, I showed experimentally that there are multiple sub-factors to speciesism. First, people believe that individuals should prioritise members of their own species over others (Species-Relativism). Second, at the same time people have a slight tendency to consider humans to be of superior value in an absolute sense (Pro-Human Bias). Third, people have a specific bias against animals, a bias people do not manifest towards non-humans with equivalent mental capacities as animals (Anti-Animal Bias). In particular the latter two factors are difficult to justify

on moral grounds and as such qualify—at least philosophically—as prejudices, which would be in line with the finding of Chapter 2 that speciesism shows psychological properties of a prejudice.

In Chapter 4, I found that people have much weaker deontological constraints against harming animals than against harming humans for the greater good. Participants considered it largely impermissible to harm humans to save many more humans, but they considered it much more permissible to harm animals to save many more animals. I explain these findings by assuming that the strength of deontological constraints is weighted by the moral status people attribute to the individual in question (Multi-level Weighted Deontology). Since species-membership itself is a crucial determinant of moral status, speciesism appears to be a key driver of deontological constraints against harm.

While all of these findings together strongly support the hypothesis that speciesism is a key driver of moral anthropocentrism, I also found that the perceived level of mental capacities of animals contributes to moral anthropocentrism. In Chapter 3, I demonstrated that participants valued animals that were described as having advanced (i.e. human-level) mental capacities more than those that were described as only having basic mental capacities. While participants always prioritised humans over animals, irrespective of mental capacity levels, they prioritised humans to a weaker extent over animals if the animals had advanced compared to basic mental capacities. This is in line with a finding of Chapter 2 that showed that participants were more willing to help a mentally severely disabled human over a chimpanzee, despite believing that the latter has more advanced mental capacities. This mental capacity effect, however, only applied when the mental capacities of *individual* animals were manipulated (Individual Mental Capacity View), not when the *typical* mental capacity level of a whole animal species was manipulated (Species-Typical Mental Capacity View). While participants valued mentally more advanced animals over mentally basic animals, they did not value mentally more advanced humans over mentally basic humans to a similar extent, suggesting that the selective application of the Individual Mental Capacity View is

itself driven by speciesism. Finally, in Chapter 4 I also found that participants had slightly stronger deontological constraints against harming beings with advanced cognitive capacities compared to beings with basic cognitive capacities. Sentience, in contrast, did not influence the strength of deontological constraints.

The Relative Strength of Speciesism and Mental Capacities

These findings make clear that there are at least two main factors driving moral anthropocentrism: speciesism and perceived mental capacities (i.e. Individual Mental Capacity View). While I did not explicitly compare the relative strength the two factors have in explaining moral anthropocentrism, I believe my studies suggest that speciesism is the stronger factor. Throughout all my studies, the species-membership manipulations generally had strong effects across various types of contexts. The effect sizes of the mental capacity manipulations, in contrast, were rather small in comparison. In the prioritisation dilemmas of Chapter 3, for example, even in the most extreme cases (e.g. pitting a human with animal-level mental capacities against a monkey with above-human-level mental capacities) participants consistently prioritised the human over the animal. Even though the mental capacity manipulations did affect people's judgments, the effects were generally small (e.g. *Figure 3.1*). In Chapter 4, this pattern was even clearer. The effect that the mental capacity manipulations had in Study 4.6 were only very small for cognitive capacities and non-existent for sentience. By contrast, the species-membership manipulations effect sizes were up to seven times larger, e.g. when comparing the effect size for the cognitive capacity manipulation in Study 4.6 with the effect size for the species-membership manipulation in Study 4.2b.

A complicating factor of determining the relative strength of speciesism and perceived mental capacities in explaining moral anthropocentrism is the fact that the former can affect the latter, which, in turn, can further enhance moral anthropocentrism. In Chapter 3, I discussed in

more detail how speciesism reduces the perceived moral status of animals both directly and indirectly via de-mentalization and a selective application of the Individual Mental Capacity View to animals but not to humans. Perhaps the key figure of this thesis, *Figure 3.7*, depicts these possible causal paths. Future research, however, needs to test and refine this framework in more detail.

Speciesism as a Form of Ingroup Favouritism

At this point, we do not have strong evidence of what the ultimate explanation of speciesism may be. But reflection on evolutionary and rational considerations suggest a preliminary, if inevitably speculative, account of the ultimate sources of speciesism.

There is extensive evidence that people categorise individuals into different groups (cf. Tajfel, Billig, Bundy, & Flament, 1971), identify with their own group (Hornsey, 2008; Tajfel, Turner, Austin, & Worchel, 1979), and prioritise members of their ingroup over members of their outgroup (Balliet, Wu, & De Dreu, 2014; Fu et al., 2012; Sherif, 1961; Yamagishi & Kiyonari, 2000; for bounded generalised reciprocity theory, cf. Yamagishi & Mifune, 2008). Ingroup favouritism is expressed in many different contexts. People have, for example, a tendency to favour others who share their ethnicity, nationality, religion, or political affiliation. (Rand et al., 2009; Whitt & Wilson, 2007). It has been argued that ingroup favouritism is an innate tendency since it can promote safety and help to encourage mutual cooperation among ingroup members (Gaertner & Insko, 2000). It seems, therefore, that there are good reasons to assume that speciesism is a form of ingroup favouritism analogous to ingroup favouritism amongst human groups.

While typical human ingroups would be far smaller than humanity itself, our similarity to other humans would be salient in contexts where a choice needs to be made between a human and a non-human. Since the differences between humans and animals are perceived as vast—in terms of biology, physical appearance, mental capacities, and behaviour—and thus the boundaries between the groups so wide and clear, one would expect ingroup favouritism between humans

and animals to be particularly strong. Indeed, research suggests that perceived similarity with outgroup members can reduce ingroup favouritism—as long as they are seen as non-threatening (Henderson-King, Henderson-King, Zhermer, Posokhova, & Chiker, 1997). Similarly, it has been shown that people have more positive reactions towards animals that are perceived as biologically, physically, mentally, or behaviourally more similar to humans than animals that are dissimilar (Burghardt & Herzog, 1989; Kellert & Berry, 1980)

Already previous research, discussed in Chapter 1, provided some evidence that the same psychological processes causing negative perceptions and treatment of human outgroup members are also active in human-animal relations. In this thesis I further support the view that people think about animals in an “us” versus “them” way. In particular, I clarified that it is primarily species-membership itself that is driving this form of ingroup favouritism. Chapter 2 showed that speciesism is psychologically similar to human-human forms of ingroup favouritism. Like racism, speciesism goes along with high social dominance orientation—a construct that is conceptually close to a generalised tendency for ingroup favouritism. The sub-factor Species-Relativism (cf. Chapter 3) appears analogous to other forms of partiality among humans, such as partiality towards one’s family members or nationalism. Like Species-Relativism, these forms of partiality show in the belief that everyone is allowed to prioritise their own group members over others. Similarly, the sub-factor Pro-Human Bias appears analogous to beliefs about the absolute superiority of certain groups over other groups. Taken together, these findings support the view that fundamentally speciesism is a form of psychological ingroup favouritism similar to forms of ingroup favouritisms amongst human groups.

Differences in Moral Thinking About Humans and Animals

Despite the fact that people value animals much less than humans, my studies suggest that moral thinking about animals does not *qualitatively* differ from moral thinking about humans.

Instead, the difference rather seems to be one of degree. There are two reasons to assume this is the case.

First, as discussed in the previous section, I provided support for the view that speciesism is a form of ingroup favouritism analogous to human-human versions of ingroup favouritism. Were moral thinking about animals fundamentally different to moral thinking about humans, we would have expected speciesism to be a psychologically distinct construct from these other types of human-human ingroup favouritism, but that was not the case.

Second, the Multi-level Weighted Deontology model presented in Chapter 4 showed that people hold weaker deontological constraints for animals than humans, but this difference is one of degree. It was not the case that people had *no* deontological constraints for animals at all, which would have suggested that different mechanisms are active when people consider whether to harm animals or humans. Instead, the findings suggest that the same psychological mechanisms are active, but that they just involve different weights, depending on whether humans or different types of animals are involved.

However, there could also be reasons that suggest that moral thinking about animals has aspects that are qualitatively different from moral thinking about humans. Evolutionarily and historically, animals have always played the roles of predators or prey for humans. It is, therefore, plausible to assume that special psychological mechanisms have evolved for thinking about animals. It has been shown, for example, that children from different cultures categorise animals and humans distinctly (Leddon et al., 2012). While my studies did not explicitly address this question, the discovery of the Anti-Animal Bias in Chapter 3 might serve as some evidence for the assumption that people think at least partially in a special way about animals.

Overall, though, my studies suggest that moral thinking about animals involves largely the same psychological attitudes and mechanisms as moral thinking about humans does. An implication of this is that psychological insights from studying how people morally value humans

and animals are in principle transferrable, as long as one accounts for the difference in degree, which can be substantial.

Limitations and Critical Reflection

Despite being able to confirm the speciesism hypothesis and explore various important aspects of speciesism, there are things I would do differently were I to start my doctoral research again.

Perhaps it would have been a better approach to have begun my research by studying the sub-factors of speciesism as well as the role mental capacities play in explaining moral anthropocentrism, as I did in Chapter 3. Instead, I began my research with the development of the Speciesism Scale because I believed that having a reliable instrument of measuring the phenomenon would facilitate studying it in greater detail later on. In retrospect, however, I believe that it would have helped to first clarify the relation between the different concepts: moral anthropocentrism, speciesism, the Mental Capacity View, and the sub-factors of speciesism. Had I started my research in this way, I would have included items into the initial item pool of the scale development that explicitly captured Species-Relativism, Pro-Human Bias, and Anti-Animal Bias. It would have been interesting to see whether factor analysis would have revealed them as separate factors. However, I expect that the extent to which these sub-factors differ between people in normal cases is very small, which means that they likely would not be detected by factor analysis but rather have to be uncovered experimentally relying on hypothetical thought experiments, as I did in Chapter 3.

A limitation of my thesis is that my empirical investigation of the sub-factors of speciesism in Chapter 3 was relatively superficial. The studies were restricted to demonstrating the existence of the different factors and did not investigate their psychological mechanisms in more detail. In particular, it remains unclear whether factors such as Species-Relativism are causal determinants

of speciesism or whether they are post-hoc rationalizations. Still, I believe that Chapter 3 makes an important contribution by giving a broad framework for moral anthropocentrism and its sub-factors, which I might investigate in more detail in future research.

In Chapter 4 I unsuccessfully tried to experimentally reduce people's individual level of speciesism with a technique that relied on text and visuals to argue for better treatment of animals. An implication of this failure was that I was not able to experimentally demonstrate that a person's individual level of speciesism contributes to their strength of deontological constraints against harming animals. Instead, I had to rely on manipulations of the species-membership and correlations between speciesism and moral judgments. An experimental manipulation of speciesism would have also proven useful to investigate its causal mechanism in more detail. Future research could try to find ways to successfully manipulate speciesism. Perhaps radically different approaches would be needed that cannot easily be conducted in an online setting. An alternative would also be to try to increase speciesism instead of reducing it. While, this arguably would not be of practical use, a robust manipulation technique could still prove useful in furthering our understanding of speciesism.

I chose to conduct most of my studies online for a number of reasons. An obvious advantage is that recruiting participants online is much faster than offline. This allowed me to conduct a large number of studies. Had I conducted all studies offline, this thesis would contain only a fraction of its current insights. Another advantage is that online participants have been shown to be more diverse (Shapiro, Chandler, & Mueller, 2013) and highly representative (Buhrmester, Kwang, & Gosling, 2011) than traditional samples (e.g. Berinsky, Huber, & Lenz, 2012; Rand, 2012)—largely because online studies allow for the recruitment of broader population samples than the often-limited university student samples. Research also suggests that data obtained via online platforms such as Amazon MechanicalTurk is of high quality (Paolacci & Chandler, 2014) and that results are comparable to results from campus studies (Bartneck,

Duenser, Moltchanova, & Zawieska, 2015). However, online studies clearly have their limitations. They are limited to asking participants questions, to which they may respond dishonestly—especially concerning sensitive moral questions. Another issue is that participants may be inattentive, which I tried to account for by including attention check questions and relying on large sample sizes whenever possible. Finally, to test whether my key findings were transferable to other samples and settings I conducted offline studies at the end of Chapter 2 and 4, which confirmed the key effects.

Future Directions

A key focus area of future research could lie on the refinement of my framework for moral anthropocentrism. In particular, the sub-factors of speciesism—Species-Relativism, Pro-Human Bias, and Anti-Animal Bias—remain poorly understood. In my research, I primarily focused on identifying the different sub-factors without studying their relative strengths and causal relations. As discussed above, it is possible, for example, that Species-Relativism is not causally driving moral anthropocentrism but rather it could be a post-hoc rationalization. Similarly, the precise causal roles of Pro-Human Bias and Anti-Animal Bias need to be studied in more depth.

It is possible that additional sub-factors of moral anthropocentrism and speciesism exist. In my studies I mainly investigated how people value large mammals like pigs, dogs or chimpanzees. Future research could expand this research by investigating how people value other, smaller animals, including small mammals, amphibia, reptiles, birds, fish, and even invertebrates such as insects. It is likely that people value smaller animals even less than larger mammals. If so, what factors explain why people value certain animals more than others. As I have shown, people value pet animals, such as dogs, more than food animals, such as pigs. This could mean that people place animals into different categories that determine their moral worth, such as food, pet, wild, pest, etc. If that is true, how do these categories relate to the Anti-Animal Bias?

One further concrete avenue for future research is to test how people perceive the moral character of others who morally value animals in a certain way. In a separate project together with collaborators, I found that people perceive speciesists (as defined by how they respond to items of the Speciesism Scale) to be similarly untrustworthy as they perceive racists and sexists to be (Everett, et al., 2018). This suggests that people are intuitively aware of the similarities between speciesism and other forms of human-human prejudice since they infer a general prejudicial attitude and, accordingly, similar character traits from people with speciesist, racist, or sexist attitudes. However, while in this project we investigated perceptions of speciesists in a general sense (as measured by the Speciesism Scale), it remains unclear how people perceive others who decide in moral dilemmas involving animals and humans in a certain way (as in Chapters 3 and 4). For example, how do people perceive a person who decides to prioritise a human over an animal compared to a person who decides to prioritise an animal over a human? Such research could potentially shed light on the ultimate causes of speciesism.

Implications

In this thesis I showed that people value animals less than humans on the basis of their species-membership, i.e. speciesism. I hope that my findings can contribute to our understanding of human morality by beginning to fill a crucial gap. Social and moral psychologists have so far paid little attention to how people value animals, despite the central role animals play in our lives.

My findings could also have implications for other disciplines, in particular moral philosophy. I hope that my findings of the different sub-factors of speciesism and their classification into Species-Relativism and Species-Absolutism will prove fruitful concepts for the philosophical discussion. This thesis and several of its hypotheses were inspired by philosophical views on the moral status of animals. The presented studies clearly confirmed some empirical hypotheses put forward by philosophers, such as the speciesism hypothesis. But my studies also

empirically refuted or refined other hypotheses philosophers have discussed. For example, the studies ruled out the hypothesis that people value humans more than animals because humans typically have advanced and animals typically basic mental capacities. My studies also refined Nozick's consideration that people might be Kantians for humans and utilitarian for animals by showing that deontological constraints against harm get weaker but do not completely disappear for animals. An empirical refutation of an implicit or explicit psychological hypothesis underpinned in a philosophical view does not necessarily imply that the view is normatively false. But these empirical findings could still inform the philosophical discussion about the moral status of animals. Ultimately, a deeper understanding of the psychology of how we value animals may even bear on the ethical evaluation of human priority over animals.

Appendix

List of Figures

Figure 2.1 Donation allocations to the charities (Study 2.4). In all three scenarios allocations significantly differed from an equal distribution of 50 donation points each. Black points represent raw data, horizontal bars represent means, rectangles represent confidence intervals, and “beans” represent smoothed densities. 44

Figure 3.1. Participants prioritised humans over monkeys in all types of dilemmas: in those where the humans had superior mental capacities, in those where humans and monkeys had equal mental capacities, and in those where monkeys had superior mental capacities. (Study 3.1) 69

Figure 3.2. Participants did not prioritise mentally superior humans over mentally inferior humans, but they did prioritise mentally superior monkeys over mentally inferior monkeys. (Study 3.1) 69

Figure 3.3. Participants tended to prioritise humans over non-humans. It made no difference whether the individual extraterrestrials (Atlans) were members of a typically mentally advanced or a typically mentally basic extraterrestrial species. The tendency to prioritise humans over extraterrestrials, however, was weaker than the tendency to prioritise humans over monkeys. For more detailed results of the individual dilemmas see *Table 3.1*. (Study 3.2) 78

Figure 3.4. Participants tended to say that humans should prioritise themselves over others, that extraterrestrials (Atlans) should sometimes prioritise members of their own species over others, and that monkeys should flip a coin to decide who to prioritise (Study 3.3a). 83

Figure 3.5. Participants tended to say that both humans and intelligent extraterrestrials should prioritise members of their own species over other individuals. In addition, there was a slight overall tendency towards prioritising humans. (Study 3.3b) 88

Figure 3.6. Participants were more likely to both not help as well as actively harm monkeys than Atlans (averaged scores). For both species, they were more willing to not help them than to actively harm them. In the harm conditions 1 stands for *Definitely harm the Atlan* and 7 for *Definitely NOT harm the Atlan* (reversed). In the help conditions 1 stands for *Definitely prioritise the human* and 7 for *Definitely prioritise the Atlan*. 4 stands for *Flip a coin to decide* in all conditions (Study 3.4). 96

Figure 3.7. A Framework for Moral Anthropocentrism. One reason people perceive animals to have low moral status because of mere species-membership, i.e. speciesism (1). Another reason is the fact that people attribute lower moral status to beings that have low perceived mental capacities, i.e. Individual Mental Capacity View (2). Furthermore, there are two indirect paths from speciesism to moral anthropocentrism. First, speciesism leads people to have lower perceived mental capacities of animals (3). Second, due to speciesism, people selectively apply the Individual Mental Capacity View only to animals but not humans (4), a phenomenon which is itself speciesist (i.e., speciesism appears to moderate the Individual Mental Capacity View). 103

Figure 4.1. Moral judgments on harming one to save five individuals, ranging from 1 (*morally wrong*), over 4 (*neither right nor wrong*), to 7 (*morally right*). Black points represent raw data, horizontal bars represent means, rectangles represent confidence intervals, and “beans” represent smoothed densities. 121

Figure 4.2. Moral judgments ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*). 126

Figure 4.3. Moral judgments of harming 10 individuals of a species to saving 100 of the same species, ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*). 131

Figure 4.4. Moral judgments of killing 10 individuals of a species to saving 100 of either the same or a different species, ranging from 1 (*absolutely morally wrong*), to 4 (*neither right nor wrong*), to 7 (*absolutely morally right*). 137

Figure 4.5. Moral judgments about harming few humans or pigs to save more, ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*). There were no differences between the speciesism-reduction and control conditions. 142

Figure 4.6. Moral judgments about harming few humans, animals, or objects to save more. Moral judgments of harming (or destroying) 10 humans, animals or objects to save 100 humans, animals or objects respectively, ranging from 1 (*morally wrong*), to 4 (*neither right nor wrong*), to 7 (*morally right*). 146

Figure 4.7. Moral judgments of harming 10 individuals of a hypothetical species to saving 100 individuals of the same species, ranging from 1 (*absolutely morally wrong*), to 4 (*neither right nor wrong*), to 7 (*absolutely morally right*). 153

Figure 4.8. Judgment about the ethical justifiability of the medical research program that harms pigs to save many more pigs or human infants to save many more human infants, ranging from 1 (*completely unjustified*), to 4 (*neither justified nor unjustified*), to 7 (*completely justified*). 159

List of Tables

Table 2.1. Factor loadings from the second EFA (Study 2.1).	24
Table 2.2. Speciesism Scale.....	28
Table 2.3. Correlations between speciesism and other constructs as well as partial correlations (pr) when controlled for SDO (Study 2.3).	38
Table 2.4. Beliefs about intelligence and suffering capability (ratings on a scale from 0 to 100) and correlations of these beliefs with speciesism (Study 2.4).....	45
Table 2.5. Hierarchical regression predicting donation allocations to “superior” individuals (humans, mentally severely disabled humans, dogs) vs. “inferior” individuals (animals, chimpanzees, pigs), displaying the standardised coefficients (beta) (Study 2.4).....	47

<i>Table 2.6.</i> Hierarchical logistic regression predicting whether participants chose the human or animal charity), displaying Exp(B) (Study 2.5).	54
<i>Table 3.1.</i> Prioritisation judgments <i>M</i> (<i>SD</i>) of Study 3.2. 1 stands for prioritising the first individual, 4 for flipping a coin to decide whom to prioritise, and 7 for prioritising the second individual. Rows represent the dilemmas in which two individuals are pitted against each other. Columns represent the type of species the respective non-human being is a member of, i.e. the three between-subjects conditions.	76
<i>Table 3.2.</i> Correlations between the three explicitly endorsed moral principles and prioritisation decisions.	89
<i>Table 3.3.</i> Means and standard deviations of the responses in the moral dilemmas involving mentally basic and mentally advanced individuals (Study 3.4). Lower scores indicate greater willingness to help humans instead of non-humans and greater willingness to actively harm non-humans to save humans (reversed).	96
<i>Table 4.1.</i> Potential moral views of when it is permissible to harm humans and animals.	113
<i>Table 4.2.</i> Moral judgments and correlations (Study 4.2).....	132
<i>Table 4.3.</i> Moral Judgments and Correlations (Study 4.3)	139
<i>Table 4.4.</i> Descriptive Statistics of Tipping Points (Study 4.3)	139
<i>Table 4.5.</i> Moral Judgment and Additional Measures (Study 4.5).....	147
<i>Table 4.6.</i> Correlations with Moral Judgment of Harming Hypothetical Individuals to Save Many Across Conditions (Study 4.6).....	154

References

- Akrami, N., Ekehammar, B., & Bergh, R. (2011). Generalized prejudice: Common and specific components. *Psychological Science*, *22*(1), 57–59.
- Allen, M. W., & 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*, *38*(2), 118–130.
- Allen, M. W., Wilson, M., Ng, S. H., & Dunne, M. (2000). Values and beliefs of vegetarians and omnivores. *The Journal of Social Psychology*, *140*(4), 405–422.
- Allport, G. W., Clark, K., & Pettigrew, T. (1954). *The nature of prejudice*.
- Altemeyer, B. (1988). *Enemies of freedom: Understanding right-wing authoritarianism*. Jossey-Bass.
- Amiot, C. E., & Bastian, B. (2015). Toward a psychology of human-animal relations. *Psychological Bulletin*, *141*(1), 6–47. <https://doi.org/10.1037/a0038147>
- Anthoine, E., Moret, L., Regnault, A., Sébille, V., & Hardouin, J.-B. (2014). Sample size used to validate a scale: a review of publications on newly-developed patient reported outcomes measures. *Health and Quality of Life Outcomes*, *12*(1), 2.
- Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., ... Rahwan, I. (2018). The Moral Machine experiment. *Nature*, *563*(7729), 59. <https://doi.org/10.1038/s41586-018-0637-6>
- Balliet, D., Wu, J., & De Dreu, C. K. W. (2014). Ingroup favoritism in cooperation: A meta-analysis. *Psychological Bulletin*, *140*(6), 1556–1581. <https://doi.org/10.1037/a0037737>
- Balter, M. (2010, April 26). Chimps Grieve Over Dead Relatives. Retrieved May 21, 2019, from Science | AAAS website: <https://www.sciencemag.org/news/2010/04/chimps-grieve-over-dead-relatives>

- Bandura, A. (1999). Moral disengagement in the perpetration of inhumanities. *Personality and Social Psychology Review*, 3(3), 193–209.
- Baron, J. (2000). *Thinking and deciding*. Cambridge University Press.
- Bartneck, C., Duenser, A., Moltchanova, E., & Zawieska, K. (2015). Comparing the similarity of responses received from studies in Amazon’s Mechanical Turk to studies conducted online and with direct recruitment. *PloS One*, 10(4), e0121595.
- Bastian, B., & Loughnan, S. (2017). Resolving the meat-paradox: A motivational account of morally troublesome behavior and its maintenance. *Personality and Social Psychology Review*, 21(3), 278–299.
- Bastian, B., Loughnan, S., Haslam, N., & Radke, H. (2012). Don’t Mind Meat? The Denial of Mind to Animals Used for Human Consumption. *Personality & Social Psychology Bulletin*, 38, 247–256. <https://doi.org/10.1177/0146167211424291>
- Bauman, C. W., McGraw, A. P., Bartels, D. M., & Warren, C. (2014). Revisiting External Validity: Concerns about Trolley Problems and Other Sacrificial Dilemmas in Moral Psychology. *Social and Personality Psychology Compass*, 8(9), 536–554. <https://doi.org/10.1111/spc3.12131>
- Bentham, J. (1780). 1988 The principles of morals and legislation. *Amherst, NY: Prometheus Books*.
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon. com’s Mechanical Turk. *Political Analysis*, 20(3), 351–368.
- Bostyn, D. H., Sevenhant, S., & Roets, A. (2018). Of mice, men, and trolleys: Hypothetical judgment versus real-life behavior in trolley-style moral dilemmas. *Psychological Science*, 29(7), 1084–1093.
- Bratanova, B., Loughnan, S., & Bastian, B. (2011). The effect of categorization as food on the perceived moral standing of animals. *Appetite*, 57(1), 193–196.

- Broom, D. M., Sena, H., & Moynihan, K. L. (2009). Pigs learn what a mirror image represents and use it to obtain information. *Animal Behaviour*, 78(5), 1037–1041.
- Brown, K., F., Simon Kroll, J., Hudson, M., Ramsay, M., Green, J., J Long, S., ... Sevdalis, N. (2010). *Factors underlying parental decisions about combination childhood vaccinations including MMR: A systematic review* (Vol. 28). <https://doi.org/10.1016/j.vaccine.2010.04.052>
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1), 3–5.
- Burghardt, G. M., & Herzog, H. A. (1989). Animals, evolution and ethics. *Perceptions of Animals in American Culture*, 129–151.
- Caviola, L., Everett, J. A. C., & Faber, N. S. (2018). The moral standing of animals: Towards a psychology of speciesism. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspp0000182>
- Cikara, M., Farnsworth, R. A., Harris, L. T., & Fiske, S. T. (2010). On the wrong side of the trolley track: Neural correlates of relative social valuation. *Social Cognitive and Affective Neuroscience*, 5(4), 404–413. <https://doi.org/10.1093/scan/nsq011>
- Cohen, C. (1986). The case for the use of animals in biomedical research. *N. Engl. J. Med.*, 315(14), 865–870.
- Cohen, P., West, S. G., & Aiken, L. S. (2014). *Applied multiple regression/correlation analysis for the behavioral sciences*. Psychology Press.
- Comrey, A. L., & Lee, H. B. (2013). *A first course in factor analysis*. Psychology press.
- Costello, K., & Hodson, G. (2010). Exploring the roots of dehumanization: The role of animal—human similarity in promoting immigrant humanization. *Group Processes & Intergroup Relations*, 13(1), 3–22.

- Costello, K., & Hodson, G. (2014). Explaining dehumanization among children: The interspecies model of prejudice. *The British Journal of Social Psychology*, *53*(1), 175–197.
<https://doi.org/10.1111/bjso.12016>
- Crimston, D., Bain, P. G., Hornsey, M. J., & Bastian, B. (2016). Moral expansiveness: Examining variability in the extension of the moral world. *Journal of Personality and Social Psychology*, *111*(4), 636–653. <https://doi.org/10.1037/pspp0000086>
- Davis, M. (1980). A Multidimensional Approach to Individual Differences in Empathy. *JSAS Catalog Sel. Doc. Psychol.*, *10*.
- Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology*, *44*(1), 113–126.
<https://doi.org/10.1037/0022-3514.44.1.113>
- DeScioli, P., & Kurzban, R. (2013). A solution to the mysteries of morality. *Psychological Bulletin*, *139*(2), 477.
- Dhont, K., & Hodson, G. (2014). Why do right-wing adherents engage in more animal exploitation and meat consumption? *Personality and Individual Differences*, *64*, 12–17.
- Dhont, K., Hodson, G., Costello, K., & MacInnis, C. C. (2014). Social dominance orientation connects prejudicial human–human and human–animal relations. *Personality and Individual Differences*, *61*, 105–108.
- Dhont, K., Hodson, G., & Leite, A. C. (2016). Common Ideological Roots of Speciesism and Generalized Ethnic Prejudice: The Social Dominance Human–Animal Relations Model (SD-HARM). *European Journal of Personality*, *30*(6), 507–522.
<https://doi.org/10.1002/per.2069>
- Diamond, C. (1978). Eating Meat and Eating People. *Philosophy*, *53*(206), 465–479.
<https://doi.org/10.1017/S0031819100026334>

- Duckitt, J., & Sibley, C. G. (2007). Right wing authoritarianism, social dominance orientation and the dimensions of generalized prejudice. *European Journal of Personality, 21*(2), 113–130.
- Ekehammar, B., & Akrami, N. (2003). The relation between personality and prejudice: a variable- and a person-centred approach. *European Journal of Personality, 17*(6), 449–464.
- Everett, J. A., Caviola, L., Savulescu, J., & Faber, N. (2018). Speciesism, generalized prejudice and perceptions of prejudiced others. *Group Processes and Intergroup Relations*.
- Everett, J. A., Faber, N. S., Savulescu, J., & Crockett, M. J. (2018). The costs of being consequentialist: Social inference from instrumental harm and impartial beneficence. *Journal of Experimental Social Psychology, 79*, 200–216.
- Everett, J. A., Pizarro, D. A., & Crockett, M. J. (2016). Inference of trustworthiness from intuitive moral judgments. *Journal of Experimental Psychology: General, 145*(6), 772.
- Everitt, B. S. (1975). Multivariate analysis: The need for data, and other problems. *The British Journal of Psychiatry, 126*(3), 237–240.
- Finnis, J. (1997). The fragile case for euthanasia: a reply to John Harris. *Euthanasia Examined: Ethical, Clinical and Legal Perspectives, 46–55*.
- Francione, G. L. (1995). *Animals Property & The Law*. Retrieved from <https://www.jstor.org/stable/j.ctt1bw1jm9>
- Fu, F., Tarnita, C. E., Christakis, N. A., Wang, L., Rand, D. G., & Nowak, M. A. (2012). Evolution of in-group favoritism. *Scientific Reports, 2*, 460.
- Gaertner, L., & Insko, C. A. (2000). Intergroup discrimination in the minimal group paradigm: Categorization, reciprocation, or fear? *Journal of Personality and Social Psychology, 79*(1), 77.
- Graça, J., Calheiros, M. M., & Oliveira, A. (2016). Situating moral disengagement: Motivated reasoning in meat consumption and substitution. *Personality and Individual Differences, 90*, 353–364.

- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, *96*(5), 1029.
- Gray, K., Young, L., & Waytz, A. (2012). Mind Perception Is the Essence of Morality. *Psychological Inquiry*, *23*(2), 101. <https://doi.org/10.1080/1047840X.2012.651387>
- Greene, J. (2009). The cognitive neuroscience of moral judgment. *The Cognitive Neurosciences*, *4*, 1–48.
- Greene, J. (2014). *Moral tribes: Emotion, reason, and the gap between us and them*. Penguin.
- Greene, J., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive Load Selectively Interferes with Utilitarian Moral Judgment. *Cognition*, *107*(3), 1144–1154. <https://doi.org/10.1016/j.cognition.2007.11.004>
- Haslam, N. (2006). Dehumanization: An Integrative Review. *Personality and Social Psychology Review*, *10*(3), 252–264. https://doi.org/10.1207/s15327957pspr1003_4
- Haslam, S. A. (1997). *Stereotyping and social influence: Foundations of stereotype consensus*.
- Henderson-King, E., Henderson-King, D., Zhermer, N., Posokhova, S., & Chiker, V. (1997). In-group favoritism and perceived similarity: A look at Russians' perceptions in the post-Soviet era. *Personality and Social Psychology Bulletin*, *23*(10), 1013–1021.
- Herek, G. M. (1998). Attitudes toward lesbians and gay men scale. *Handbook of Sexuality-Related Measures*, 392–394.
- Herzog, H. A., & Mathews, S. (1997). Personality and attitudes toward the treatment of animals. *Society & Animals*, *5*(2), 169–175.
- Herzog, H., Betchart, N. S., & Pittman, R. B. (1991). Gender, sex role orientation, and attitudes toward animals. *Anthrozoös*, *4*(3), 184–191.
- Hewstone, M., Stroebe, W., & Jonas, K. (2016). *An Introduction to Social Psychology*. John Wiley & Sons.

- Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin*, *142*(11), 1179.
- Hornsey, M. J. (2008). Social Identity Theory and Self-categorization Theory: A Historical Review. *Social and Personality Psychology Compass*, *2*(1), 204–222.
<https://doi.org/10.1111/j.1751-9004.2007.00066.x>
- Horta, O. (2010). What is Speciesism? *Journal of Agricultural and Environmental Ethics*, *23*(3), 243–266. <https://doi.org/10.1007/s10806-009-9205-2>
- Hu, L., & Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal*, *6*(1), 1–55.
- Huber, S., & Huber, O. W. (2012). The centrality of religiosity scale (CRS). *Religions*, *3*(3), 710–724.
- Jackson, D. L. (2001). Sample size and number of parameter estimates in maximum likelihood confirmatory factor analysis: A Monte Carlo investigation. *Structural Equation Modeling*, *8*(2), 205–223.
- Joy, M. (2011). *Why we love dogs, eat pigs, and wear cows: An introduction to carnism*. Conari press.
- Kagan, S. (2016). What’s Wrong with Speciesism? (Society for Applied Philosophy Annual Lecture 2015). *Journal of Applied Philosophy*, *33*(1), 1–21.
<https://doi.org/10.1111/japp.12164>
- Kahane, G., Everett, J. A. C., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review*, *125*(2), 131–164. <https://doi.org/10.1037/rev0000093>
- Kahane, G., Everett, J. A. C., Earp, B. D., Farias, M., & Savulescu, J. (2015). “Utilitarian” judgments in sacrificial moral dilemmas do not reflect impartial concern for the greater good. *Cognition*, *134*, 193–209. <https://doi.org/10.1016/j.cognition.2014.10.005>

- Kant, I. (1785). *Groundwork of the Metaphysics of Morals*.
- Kasperbauer, T. J. (2017). *Subhuman: The Moral Psychology of Human Attitudes to Animals*. Oxford University Press.
- Kay, A. C., & Jost, J. T. (2003). Complementary justice: effects of "poor but happy" and "poor but honest" stereotype exemplars on system justification and implicit activation of the justice motive. *Journal of Personality and Social Psychology*, 85(5), 823.
- Keller, S. (2013). *Partiality* (Vol. 40). Princeton University Press.
- Kellert, S. R., & Berry, J. K. (1980). *Knowledge, Affection and Basic Attitudes Toward Animals in American Society. Phase III*.
- Kilkenny, C., Browne, W. J., Cuthill, I. C., Emerson, M., & Altman, D. G. (2010). Improving bioscience research reporting: The ARRIVE guidelines for reporting animal research. *PLoS Biology*, 8(6), e1000412.
- Killoren, D., & Streiffer, R. (2019). Utilitarianism about animals and the moral significance of use. *Philosophical Studies*, 1–21.
- Kim, R., Kleiman-Weiner, M., Abeliuk, A., Awad, E., Dsouza, S., Tenenbaum, J., & Rahwan, I. (2018). A Computational Model of Commonsense Moral Decision Making. *ArXiv:1801.04346 [Cs]*. Retrieved from <http://arxiv.org/abs/1801.04346>
- Kozak, M. N., Marsh, A. A., & Wegner, D. M. (2006). What do i think you're doing? Action identification and mind attribution. *Journal of Personality and Social Psychology*, 90(4), 543.
- Leddon, E. M., Waxman, S. R., Medin, D. L., Bang, M., Washinawatok, K., & Hayes, M. (2012). One animal among many? Children's understanding of the relation between humans and nonhuman animals. *Psychology of Culture. In Psychology of Emotions, Motivations and Actions: Focus on Civilizations and Cultures Series*, 105–126.

- Loughnan, S., Haslam, N., & Bastian, B. (2010). The role of meat consumption in the denial of moral status and mind to meat animals. *Appetite*, *55*(1), 156–159.
<https://doi.org/10.1016/j.appet.2010.05.043>
- Low, P., Panksepp, J., Reiss, D., Edelman, D., Van Swinderen, B., Low, P., & Koch, C. (2012). Cambridge Declaration on Consciousness in Non-Human Animals. *Francis Crick Memorial Conference on Consciousness in Human and Non-Human Animals*. University of Cambridge, Cambridge.
- Marsh, H. W., Balla, J. R., & McDonald, R. P. (1988). Goodness-of-fit indexes in confirmatory factor analysis: The effect of sample size. *Psychological Bulletin*, *103*(3), 391.
- McConahay, J. B. (1986). *Modern racism, ambivalence, and the modern racism scale*. Retrieved from <http://doi.apa.org/psycinfo/1986-98698-004>
- McFarland, S. (2010). Authoritarianism, social dominance, and other roots of generalized prejudice. *Political Psychology*, *31*(3), 453–477.
- McMahan, J. (2002). *The ethics of killing: Problems at the margins of life*. Oxford University Press, USA.
- Mendl, M., Held, S., & Byrne, R. W. (2010). Pig cognition. *Current Biology*, *20*(18), R796–R798.
- Mikhail, J. (2007). Universal moral grammar: Theory, evidence and the future. *Trends in Cognitive Sciences*, *11*(4), 143–152.
- Mill, J. S. (1861). Considerations on Representative Democracy. *Parker, Son and Bourn, London*.
- Miller, D. (2005). Reasonable partiality towards compatriots. *Ethical Theory and Moral Practice*, *8*(1–2), 63–81.
- Monteiro, C. A., Pfeiler, T. M., Patterson, M. D., & Milburn, M. A. (2017). The Carnism Inventory: Measuring the ideology of eating animals. *Appetite*, *113*, 51–62.
- Mundfrom, D. J., Shaw, D. G., & Ke, T. L. (2005). Minimum sample size recommendations for conducting factor analyses. *International Journal of Testing*, *5*(2), 159–168.
- Nichols, S., & Mallon, R. (2006). Moral dilemmas and moral rules. *Cognition*, *100*(3), 530–542.

- Nijssen, S. R., Müller, B. C., Baaren, R. B. van, & Paulus, M. (2019). Saving the Robot or the Human? Robots Who Feel Deserve Moral Care. *Social Cognition, 37*(1), 41–52.
- Nozick, R. (1974). Constraints and Animals. In *Anarchy, State and Utopia* (pp. 35–42). New York: Basic Books.
- Nuffield Council on Bioethics (Ed.). (2005). *The ethics of research involving animals*. London: Nuffield Council on Bioethics.
- Nunnally, J. C., Bernstein, I. H., & Berge, J. M. ten. (1967). *Psychometric theory* (Vol. 226). McGraw-hill New York.
- Nussbaum, M. C. (1995). Objectification. *Philosophy & Public Affairs, 24*(4), 249–291.
- Osborne, J. W., Costello, A. B., & Kellow, J. T. (2008). Best practices in exploratory factor analysis. *Best Practices in Quantitative Methods, 86–99*.
- Paolacci, G., & Chandler, J. (2014). Inside the Turk: Understanding Mechanical Turk as a participant pool. *Current Directions in Psychological Science, 23*(3), 184–188.
- Penner, L. A., Fritzsche, B. A., Craiger, J. P., & Freifeld, T. R. (1995). Measuring the prosocial personality. *Advances in Personality Assessment, 10*, 147–163.
- Petrinovich, L., O’neill, P., & Jorgensen, M. (1993). An empirical study of moral intuitions: Toward an evolutionary ethics. *Journal of Personality and Social Psychology, 64*(3), 467.
- Piazza, J., & Loughnan, S. (2016). When meat gets personal, animals’ minds matter less: Motivated use of intelligence information in judgments of moral standing. *Social Psychological and Personality Science, 7*(8), 867–874.
<https://doi.org/10.1177/19485506166660159>
- Piazza, J., Ruby, M. B., Loughnan, S., Luong, M., Kulik, J., Watkins, H. M., & Seigerman, M. (2015). Rationalizing meat consumption. The 4Ns. *Appetite, 91*, 114–128.
<https://doi.org/10.1016/j.appet.2015.04.011>

- Plous, S. (1993). Psychological mechanisms in the human use of animals. *Journal of Social Issues*, 49(1), 11–52.
- Poteat, V. P., & Mereish, E. H. (2012). Ideology, prejudice, and attitudes toward sexual minority social policies and organizations. *Political Psychology*, 33(2), 211–224.
- Pratto, F., Sidanius, J., Stallworth, L. M., & Malle, B. F. (1994). Social dominance orientation: A personality variable predicting social and political attitudes. *Journal of Personality and Social Psychology*, 67(4), 741.
- Pugh, J., Kahane, G., & Savulescu, J. (2013). Cohen’s conservatism and human enhancement. *The Journal of Ethics*, 17(4), 331–354.
- Quinn, W. (1984). Abortion: identity and loss. *Philos. Public Aff.*, 13(1), 24–54.
- Rand, D. G. (2012). The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, 299, 172–179.
- Rand, D. G., Pfeiffer, T., Dreber, A., Sheketoff, R. W., Wernerfelt, N. C., & Benkler, Y. (2009). Dynamic remodeling of in-group bias during the 2008 presidential election. *Proceedings of the National Academy of Sciences*, 106(15), 6187–6191.
- Regan, T. (1987). The case for animal rights. In *Advances in animal welfare science 1986/87* (pp. 179–189). Springer.
- Regan, T., & Singer, P. (1989). *Animal rights and human obligations*.
- Rothgerber, H. (2013). Real men don’t eat (vegetable) quiche: Masculinity and the justification of meat consumption. *Psychology of Men & Masculinity*, 14(4), 363.
- Rozin, P., Hormes, J. M., Faith, M. S., & Wansink, B. (2012). Is meat male? A quantitative multimethod framework to establish metaphoric relationships. *Journal of Consumer Research*, 39(3), 629–643.
- Ruby, M. B., & Heine, S. J. (2011). Meat, morals, and masculinity. *Appetite*, 56(2), 447–450.

- Ryder, R. D. (2006). Speciesism in the Laboratory. In *In defense of animals: The second wave* (pp. 87–103). Malden: Blackwell Publishing.
- Ryder, R. D. (2010). Speciesism again: The original leaflet. *Critical Society*, 2, 1–2.
- Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.
- Shapiro, D. N., Chandler, J., & Mueller, P. A. (2013). Using Mechanical Turk to study clinical populations. *Clinical Psychological Science*, 1(2), 213–220.
- Sherif, M. (1961). *Intergroup conflict and cooperation: The Robbers Cave experiment* (Vol. 10). University Book Exchange Norman, OK.
- Singer, P. (1975). *Animal liberation: A new ethic for our treatment of animals*. New York: HarperCollins.
- Singer, P. (1979). *Practical Ethics* (2 edition). Cambridge ; New York: Cambridge University Press.
- Singer, P. (1981). *The Expanding Circle: Ethics, Evolution, and Moral Progress*. Retrieved from <https://www.jstor.org/stable/j.ctt7sg4n>
- Soper, D. (2017). Statistics Calculator. Retrieved July 27, 2018, from <https://www.danielsoper.com/statcalc/default.aspx>
- Stanovich, K. E., & West, R. F. (1997). Reasoning independently of prior belief and individual differences in actively open-minded thinking. *Journal of Educational Psychology*, 89(2), 342.
- Starmans, C., & Friedman, O. (2016). If I am free, you can't own me: Autonomy makes entities less ownable. *Cognition*, 148, 145–153. <https://doi.org/10.1016/j.cognition.2015.11.001>
- Steiger, J. H. (2007). Understanding the limitations of global fit assessment in structural equation modeling. *Personality and Individual Differences*, 42(5), 893–898.
- Stern, C., West, T. V., Jost, J. T., & Rule, N. O. (2013). The politics of gaydar: Ideological differences in the use of gendered cues in categorizing sexual orientation. *Journal of Personality and Social Psychology*, 104(3), 520.
- Swim, J. K., Aikin, K. J., Hall, W. S., & Hunter, B. A. (1995). Sexism and racism: Old-fashioned and modern prejudices. *Journal of Personality and Social Psychology*, 68(2), 199.

- Sytsma, J., & Machery, E. (2012). The Two Sources of Moral Standing. *Review of Philosophy and Psychology*, 3(3), 303–324. <https://doi.org/10.1007/s13164-012-0102-7>
- Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology*, 1(2), 149–178.
- Tajfel, H., Turner, J. C., Austin, W. G., & Worchel, S. (1979). An integrative theory of intergroup conflict. *Organizational Identity: A Reader*, 56–65.
- Taylor, N., & Signal, T. D. (2005). Empathy and attitudes to animals. *Anthrozoös*, 18(1), 18–27.
- Thomson, J. J. (1984). The trolley problem. *Yale LJ*, 94, 1395.
- Tooby, J., & Cosmides, L. (2010). Groups in mind: The coalitional roots of war and morality. *Human Morality and Sociality: Evolutionary and Comparative Perspectives*, 91–234.
- Topolski, R., Weaver, J. N., Martin, Z., & McCoy, J. (2013). Choosing between the emotional dog and the rational pal: A moral dilemma with a tail. *Anthrozoös*, 26(2), 253–263.
- Varner, G. E. (1994). The prospects for consensus and convergence in the animal rights debate. *Hastings Center Report*, 24(1), 24–28.
- Waytz, A., Gray, K., Epley, N., & Wegner, D. M. (2010). Causes and consequences of mind perception. *Trends in Cognitive Sciences*, 14(8), 383–388.
- Whitt, S., & Wilson, R. K. (2007). The dictator game, fairness and ethnicity in postwar Bosnia. *American Journal of Political Science*, 51(3), 655–668.
- Williams, B. (2009). The human prejudice. In 3: Vol. 77. *Peter Singer Under Fire: The Moral Iconoclast Faces His Critics*. Open Court.
- Williams, M., & Eberhardt, J. (2008). Biological conceptions of race and the motivation to cross racial boundaries. *Journal of Personality and Social Psychology*, 94(6), 1033.
- Wuensch, K. L., Jenkins, K. W., & Poteat, G. M. (2002). Misanthropy, idealism and attitudes towards animals. *Anthrozoös*, 15(2), 139–149.

Yamagishi, T., & Kiyonari, T. (2000). The Group as the Container of Generalized Reciprocity.

Social Psychology Quarterly, 63(2), 116–132. <https://doi.org/10.2307/2695887>

Yamagishi, T., & Mifune, N. (2008). Does shared group membership promote altruism? Fear,

greed, and reputation. *Rationality and Society*, 20(1), 5–30.

Yuker, H., Block, J., & Young, J. (1970). The measurement of attitudes towards disabled persons

scale. *Albertson, NY: Ina Mend Institution.*