

Utilitarianism about animals and the moral significance of use

David Killoren¹  · Robert Streiffer²

Published online: 18 December 2018
© Springer Nature B.V. 2018

Abstract The Hybrid View endorses utilitarianism about animals and rejects utilitarianism about humans. This view has received relatively little sustained attention in the philosophical literature. Yet, as we show, the Hybrid View underlies many widely held beliefs about zoos, pet ownership, scientific research on animal and human subjects, and agriculture. We develop the Hybrid View in rigorous detail and extract several of its main commitments. Then we examine the Hybrid View in relation to the view that human use of animals constitutes a special relationship. We show that it is intuitively plausible that our use of animals alters our moral obligations to animals. That idea is widely believed to be incompatible with the sort of utilitarian approach in animal ethics that is prescribed by the Hybrid View. To overturn that conventional wisdom, we develop two different principles concerning the moral significance of human use of animals, which we call the Partiality Principle and the Strengthening Principle. We show that the Partiality Principle is consistent with several key commitments of the Hybrid View. And, strikingly, we show that the Strengthening Principle is fully consistent with all of the main commitments of the Hybrid View. Thus we establish the surprising result that utilitarians about animals can coherently offer a robust and intuitively appealing account of the moral significance of animal use.

Keywords Utilitarianism · Consequentialism · Animal ethics · Special relationships · Use

✉ David Killoren
david.j.killoren@gmail.com

Robert Streiffer
rstreiffer@wisc.edu

¹ Australian Catholic University, Melbourne, Australia

² University of Wisconsin-Madison, Madison, USA

1 Introduction

On one view, utilitarianism is mistaken about how we ought to treat human beings, but is more or less correct about non-human animals. Robert Nozick considers (without endorsing) a version of this view, which he calls *utilitarianism for animals, Kantianism for people* (1974: 39). Nozick's name for the view is overly narrow (not all non-utilitarian approaches are Kantian). Here's the name we'll use: *the Hybrid View*.

All versions of the Hybrid View (as we define it) endorse utilitarianism about animals and reject utilitarianism about humans. We explain what it means to be a utilitarian about animals in Sect. 3. But before we do that, we show in Sect. 2 that the Hybrid View underlies common views about animal research, agriculture, and zoos. The Hybrid View has been mostly overlooked in the philosophical literature, but is at least tacitly endorsed by a great many non-philosophers, including professionals and policy-makers in science and other areas where ethical concerns about animals are highly practically relevant. For this reason, sustained philosophical attention to the Hybrid View is overdue.

Among recent views in philosophical ethics, Jeff McMahan's Two-Tiered Account (2002: 245–251) comes closest to the Hybrid View [n.b., McMahan professes agnosticism on key features of the Two-Tiered Account (2002: 260)]. According to the Two-Tiered Account, the *morality of respect* governs in the realm of persons (including most but not all humans) while the *morality of interests* governs in the realm of non-persons (including most but not all animals).

McMahan's Two-Tiered Account does not count as a version of the Hybrid View for two reasons. First, while the morality of interests eschews certain non-utilitarian elements [e.g., a requirement to respect dignity (2002: 260)], it is not an unrestrictedly utilitarian principle. Second, the Two-Tiered Account distinguishes between persons and non-persons, which is different from the Hybrid View's distinction between humans and animals. Nevertheless, the Two-Tiered Account and the Hybrid View belong to the same family: both posit that different classes of individuals belong to different moral realms and that different principles ought to govern our treatment of individuals in those different classes.¹

In developing the Two-Tiered Account, McMahan fails to satisfactorily address a difficulty for theories of this family: what to do when our actions affect individuals in *both* classes, e.g., when our actions benefit humans but harm animals. Given that such cases are rife, there is a pressing need to clarify what it even means to posit the sort of moral separation that is characteristic of both the Two-Tiered Account and the Hybrid View. In Sect. 3, we'll address that difficulty and uncover core commitments of the Hybrid View.

Our aim is not to defend the Hybrid View. Rather, we aim to rigorously characterize the Hybrid View and to consider the extent to which it can be reconciled with the view that humans stand in special relationships to animals.

¹ For other theories within this family, see (Cohen 1986; Cohen and Regan 2001; Singer 2011: 71–93; Varner 2012; Regan 1983: 246).

We're especially interested in the relationship between an animal and a human who *uses* that animal for human purposes.

In Sect. 4, we explain why proponents of the Hybrid View should want to accommodate the view that human use of animals gives rise to special relationships. We identify two ways of understanding the moral significance of use—the Partiality Principle (Sect. 5) and the Strengthening Principle (Sects. 6, 7)—and explore their compatibility with the Hybrid View and its associated commitments.

As will be discussed below, both principles look like deontological principles incompatible with a utilitarian approach in animal ethics. Yet we'll show that the Strengthening Principle can be incorporated into a version of the Hybrid View, allowing the Hybrid View to be more accommodating of characteristically deontological claims than it initially appears to be. Further, we show that the Partiality Principle can be incorporated into a theory that shares some of the same utilitarian attractions as the Hybrid View, although it falls short of being a version of the Hybrid View.

Our results in this paper are significant not only for animal ethics but also for ethical theory generally. Many philosophers believe that utilitarianism is incompatible with the possibility that special relationships matter morally in and of themselves. W.D. Ross neatly expresses that idea: “If the only duty is to produce the maximum of good, the question who is to have the good—whether it is myself, or my benefactor, or a person to whom I have made a promise to confer that good on him, or a mere fellow man to whom I stand in no such special relation—should make no difference to my having a duty to produce that good” (Ross 1930: 22). And this idea remains widely accepted today (Crisp m.s.). We provide a novel argument against that conventional wisdom.

2 The Hybrid View underlies common views about zoos, research, and agriculture

To date, there has been little psychological research on whether folk intuitions involve a version of the Hybrid View.² But common views about and regulations governing institutions involving the use of animals provide ample evidence that many are inclined to some form of the Hybrid View.

Let's start with zoos. In some zoos, animals are treated with cruelty and neglect. But in well-run zoos, animals are treated well and are, on the whole, happy. Many are skeptical that there are grounds for morally objecting to well-run zoos.

To be sure, some believe that even the best zoos are unethical. Lori Gruen says that confining animals in zoos “is an exercise of domination, and it violates [the animals'] Wild dignity, even if it doesn't cause any obvious suffering” (2011: 155). Here Gruen suggests the non-utilitarian idea that features of an animal's captivity can matter independently of their effects on the animal's welfare. But Gruen's views

² Lucius Caviola et al. (in preparation) have obtained survey evidence that people are more willing to harm a few animals to save many animals than to harm a few humans to save many humans, which suggests that people hold a utilitarian view about animals but a deontological view about humans.

are not the norm: vast crowds of ordinary people who consider themselves to be morally decent routinely enjoy visiting zoos they believe to be well-run. This suggests that many believe that as long as zoos ensure that their captives are sufficiently happy, there is no basis for moral objection, which in turn suggests that many people have utilitarian intuitions about zoos.

By contrast, when we consider the idea of human zoos, purely utilitarian intuitions become scarce among decent people. It seems that there is something seriously wrong with confining humans in zoos, even in a hypothetical case in which the captive humans are treated exceptionally well and are very happy.

And one finds a similar combination of views in many other human-animal relationships that involve captivity. For example, most people have no principled moral objection to responsible ownership of non-human pets, yet ownership of human pets seems abhorrent in principle. These views can be elegantly explained in terms of the Hybrid View.

Comparing human research regulations with animal research regulations also suggests a commitment to a Hybrid View, albeit one implemented through historical accidents and political compromises. Human subjects research regulations include a variety of deontological elements. For example, in the US and other jurisdictions, there is a near absolute prohibition on using children in non-therapeutic research that is more than a minor increment over minimal risk (Ross 2005; Gennet and Altavilla 2016), even though such research could be beneficial in the aggregate. This prohibition is largely motivated by deontological concern about harming people without their consent.

Yet, with the exception of a widely publicized recent NIH funding decision about chimpanzees, no remotely comparable prohibition exists in US regulations on animal research. This reflects a broadly utilitarian approach that allows for the use of animals in significantly harmful, non-therapeutic research despite animals' inability to consent. As Gary Varner observes, most animal researchers agree with Peter Singer's utilitarian view of animals, even as they disagree with Singer about whether that view supports or condemns wide swaths of animal research (Varner 1994: 26). The few elements of US regulations on animal research that take a non-utilitarian form, such as the general prohibition on using paralytics without anesthesia, would seem to be justified as wise policy from a purely utilitarian perspective. Even the recent US decision regarding chimpanzees was largely driven by the empirical finding that "most current use of chimpanzees for biomedical research is unnecessary" (Institute of Medicine 2011: 66–67).

The Nuffield Council on Bioethics found that a utilitarian cost–benefit analysis serves as "the cornerstone" of animal research regulations in the UK (Nuffield 2005: 27, 52). Admittedly, the Council shies away from classifying the regulations as endorsing a utilitarian view of animals, noting that the regulations include a de facto ban on research involving great apes, a prohibition on harmful experiments for entertainment, a prohibition on using animals in cosmetics research, and a prohibition on procedures likely to cause unalleviated severe pain or distress (Nuffield 2005: 52, 53). Yet such prohibitions at the level of policy can also be given a utilitarian justification. Regarding the EU more generally, Andrew Knight

goes so far as to say that a “utilitarian cost:benefit analysis underpins all fundamental regulations governing animal experimentation” (2011: 3, 4).

The Hybrid View also seems implicit in widespread views about agriculture. Consider the view that Jeff McMahan calls *humane omnivorism*, according to which “factory farming is objectionable because of the suffering it inflicts on animals, [but] it is permissible to eat animals if they are reared humanely and killed with little or no pain or terror” (2016: 65) [c.f. “benign carnivorism” in McMahan (2008)]. The well-known agricultural industry consultant Temple Grandin is a devotee of this view: she has devised and advocated for handling and slaughter methods designed to reduce the suffering of farm animals, so she evidently thinks that the happiness of animals matters. Yet she still endorses the slaughter that is essentially involved in animal agriculture. Grandin’s humane omnivorism is thus consistent with (though may not entail) a utilitarian approach to animal ethics. Relatedly, Peter Sandøe observes that farm animal legislation often “has strong affinities to a utilitarian way of thinking” (2003: 473). By contrast, few would defend the use of humans for food even if it could somehow be justified on utilitarian grounds.

Some aspects of common thinking about animal ethics are at odds with the Hybrid View.³ Nevertheless, the points we have covered in this section establish that the Hybrid View captures a significant fraction of common beliefs and practices regarding humans and animals.⁴ For this reason the Hybrid View is deserving of serious attention from philosophers. We turn now to the task of giving a rigorous characterization of the Hybrid View’s core commitments.

3 The Hybrid View

What does it even mean to be a utilitarian about animals but not about humans?

A first difficulty in answering this question derives from the fact that utilitarianism is a diverse family of theories. We’ll want to understand the Hybrid View in a way that is neutral between the major branches of the utilitarian family (though we’ll set aside rule-utilitarian and expected-utility variants). So we need to identify commitments that utilitarians have in common. For starters, we suggest that utilitarians of all major types will agree to the following consequentialist thesis:

³ Consider a scene from the early-90s movie *Fierce Creatures*: Willa Weston objects to Vince McCain’s attempts to raise zoo revenue by securing sponsorships for the animals because it is “degrading to the animals.” Pointing to a sign on a tiger which has a picture of a bottle of Absolut vodka next to the slogan “Absolut Fierce,” she says “That is unacceptable!” Willa’s concern about degradation seems not to derive from concern about welfare and thus seems to be at odds with utilitarianism about animals. Insofar as Willa’s concern would be shared by many ordinary people, it illustrates that not all judgments about animals are in line with the Hybrid View. Thanks to an anonymous referee for pushing us on this point.

⁴ There is, of course, a further question as to why people who hold the Hybrid View do so. It may be an expression of speciesism—an irrational bias toward members of one’s own species—or it may be a consequence of views about the different kinds of values that can be realized in the lives of animals versus the lives of humans, values that call for different kinds of moral responses. We needn’t take a stand on this question here.

Optimific Is Always Okay: It is never wrong, and hence it is always permissible, for an agent to act optimifically (where an action is optimific iff its consequences are at least as good as the consequences of any alternative action open to the agent).⁵

To see why *Optimific Is Always Okay* will be ecumenically appealing among utilitarians, consider utilitarianism's three major varieties: maximizing, satisficing, and scalar.

Maximizing utilitarians hold that we always and everywhere have an (all things considered) obligation to act optimifically (Moore 1903: 76, 77, 197, 198; Brandt 1988: 342; Mogensen 2016: 215). Those who accept such a view will endorse *Optimific Is Always Okay*.

Satisficing utilitarians hold that we are always and everywhere obligated to do what is *at least good enough*, and that in some cases an action may be good enough without being optimific (Slote and Pettit 1984; Jamieson and Elliot 2009). For the satisficer, optimific actions are in many cases supererogatory, but will always be above the "good enough" threshold and so will never be impermissible.

Scalar utilitarians do away with deontic properties altogether and thus do not regard any actions as obligatory, forbidden, right, wrong, etc. (Norcross 2006). For the scalar utilitarian, optimific actions are never wrong because *no* actions are wrong.⁶

By contrast, deontological theories characteristically imply that it is sometimes wrong to act optimifically (Kagan 1989; Kamm 2001: 207–289, 2007: 11–47, 130–189; Mogensen 2016; Sinclair 2017). For example, a standard deontological claim, anathema to utilitarians, is that it can be morally wrong to kill one person in order to harvest her organs to save five even when doing so is optimific. Thus, *Optimific Is Always Okay* represents a distinction between deontologists and utilitarians.

Given that utilitarianism is committed to *Optimific Is Always Okay*, the Hybrid View should involve some restricted version of *Optimific Is Always Okay* that applies narrowly within the realm of animals and does not apply in the realm of humans. But it is not immediately obvious how to construct such a restricted version.

As a first pass, one might think that *Optimific Is Always Okay* applied narrowly to animals would look like this:

The Animal-Affecting Criterion: If ϕ -ing affects animals, then ϕ -ing is morally permissible if ϕ -ing is optimific.

But this principle is unsatisfactory for the simple reason that many actions affect animals and humans. To see why this creates a problem, consider the following

⁵ For consequentialists, the goodness of an act's consequences can include any intrinsic value the act itself might have (Thomson 1994: 14).

⁶ Recently, some scalar consequentialists (Sinhababu 2018) have argued that actions have deontic properties but that those properties are matters of degree. These theorists too will typically sign on for *Optimific Is Always Okay*.

variant of the footbridge trolley case. A runaway trolley is headed toward five innocent people. You can push a large man from a footbridge into the trolley's path, killing the man, but stopping the trolley and sparing the five. Here's the twist: if the large man is pushed into the trolley's path, his large body will provide a hearty feast for vultures; but the five are so small that if the trolley is allowed to kill them, the vultures will be left with only a light snack.

A standard deontological view is that it is wrong to push the large man. So the Hybrid View should be consistent with that view, since the Hybrid View is non-utilitarian about humans. But the Animal-Affecting Criterion implies that it is morally permissible to push the large man, simply because the vultures will be affected by this decision and pushing the large man is optimific. Given this, the Animal-Affecting Criterion cannot be built into the Hybrid View.

Here is a second principle that might be thought to represent Optimific Is Always Okay applied within the realm of animals:

The Not-Human-Affecting Criterion: If ϕ -ing does not affect human beings, then it is morally permissible to ϕ if ϕ -ing is optimific.

This principle seems to cohere with the sorts of views that motivate the Hybrid View, but it is largely uninformative. Nearly all actions we have an interest in ethically evaluating will have *some* effects on humans, so the antecedent will rarely be satisfied in real-world cases.

Here is a third possibility:

The Moral Significance Criterion: If ϕ -ing affects animals in a morally significant way but does not affect humans in any morally significant way, then ϕ -ing is morally permissible if ϕ -ing is optimific.

Unlike the Not-Human-Affecting Criterion, the Moral Significance Criterion has substantial implications for real-world cases. For there are many cases in which an action has morally significant effects on animals but no morally significant effects on humans.

Nevertheless, the Moral Significance Criterion does not fully capture the Hybrid View, for it does not address actions that have morally significant effects on both humans and animals. Thus, the Moral Significance Criterion does not capture many of the sorts of views we discussed in Sect. 2 to motivate the Hybrid View. For example, the Moral Significance Criterion cannot be invoked to explain the alleged permissibility of harmful animal research performed for the sake of greater human benefit. This means that although the Moral Significance Criterion should be regarded as a commitment of the Hybrid View, it cannot be taken to be exhaustive of the Hybrid View's commitments.

Here is a criterion that captures the sorts of views we've just mentioned:

The Negative Criterion: If ϕ -ing affects both humans and animals in morally significant ways and ϕ -ing is optimific, then ϕ -ing is morally wrong only if ϕ -ing is morally wrong in virtue of a relation between ϕ -ing and those humans who will be affected by ϕ -ing.

We believe that those who want to endorse the Hybrid View will sign on for the Negative Criterion. And we think that a commitment to the Negative Criterion is a distinctive (though perhaps not unique) feature of the Hybrid View.

Consider a case in which animals are harmed in research that will provide benefits for humans and thus affects animals and humans in morally significant ways. Suppose that the research is optimific (because the benefits exceed the harms and there is no better alternative). Then the Negative Criterion implies that the research is not wrong *unless* it is wrong in virtue of some relation between the research and the humans who will be affected by it. This seems clearly to be a judgment that defenders of the Hybrid View as discussed in Sect. 2 will want to endorse.

In light of these considerations, the Hybrid View should be interpreted as including a commitment to both the Moral Significance Criterion and the Negative Criterion.

Thus far we've been considering features that utilitarianism shares with other consequentialist theories: Optimific Is Always Okay is a consequentialist principle, not a specifically utilitarian one. Now we need to consider some distinguishing features of utilitarianism in particular.

We take utilitarianism (as understood in modern philosophical vernacular) to involve two major commitments. First, utilitarians are *welfarists about value*: utilitarians believe that only states that constitute welfare are intrinsically (non-instrumentally) good or bad. Second, utilitarians are *subjectivists about welfare*: utilitarians believe that mental states either constitute welfare or determine which states constitute welfare.

There are three main theories of welfare: hedonism, preferentism, and objective-list theories (Parfit 1984: 493–502; Kagan 1998: 29–41; c.f. Woodard 2013). According to hedonism, only hedonic states constitute welfare. According to preferentism, only desire satisfaction and desire frustration constitute welfare. According to objective-list theories, other states (beyond hedonic states and desire satisfaction and frustration) constitute welfare. Hedonism and preferentism imply that mental states play a fundamental role in determining welfare: for the hedonist, hedonic states constitute welfare; for the preferentist, desires determine which further states constitute welfare. So these are subjectivist theories of welfare, in contrast to objective-list theories.

Welfarism about value is compatible with any substantive theory of welfare. But, to our knowledge, all utilitarians today endorse some form of subjectivism about welfare and no utilitarians accept any objective-list theory.⁷ So we think it

⁷ Classical utilitarians are hedonists: see, e.g., Bentham (1789), Mill (1861), Sidgwick (1907) and De Lazari-Radek and Singer (2014). Post-classical utilitarians who aren't hedonists are generally preferentists: see, e.g., Smart and Williams (1973: 79–80), Hare (1981), Singer (2011: 71–93) and c.f. Railton (1984). Historically, some philosophers who identified as utilitarians held that states other than welfare are intrinsically valuable. G.E. Moore's "ideal utilitarianism" implies, for example, that beauty can be intrinsically good even if it is not enjoyed by anyone (Moore 1903: 135–136). But this reflects a usage of 'utilitarianism' that predates the modern distinction between consequentialism and utilitarianism.

reasonable to construe utilitarianism in the modern context as committed to subjectivism about welfare.

Let us now construct versions of these commitments that are restricted specifically to animals. This will require an intuitive distinction between animal states and other sorts of states. Here are some examples of animal states: a given animal has a headache; a given animal wants food; a given animal is confined in a zoo; a given animal is treated in an undignified way. By contrast, the beautiful symmetry of a painting hanging in the Louvre, or the dignity of a courageously suffering human being, are not animal states.

We suggest that utilitarians about animals are committed to a view we'll call *welfarism about animal value*: the view that animal states are intrinsically good (bad) only if they constitute positive (negative) welfare. This view leaves open the possibility that a painting's symmetry or a human being's dignity can be intrinsically good even if they do not constitute anyone's welfare. But the view implies that, e.g., an animal's freedom or dignity cannot be intrinsically good unless they constitute welfare.⁸

We further suggest that utilitarians about animals are *subjectivists about animal welfare*: they hold that *either* (1) an animal's welfare is constituted by (some of) the animal's mental states or (2) the animal's mental states determine which further states constitute the animal's welfare. The conjunction of welfarism about animal value and subjectivism about animal welfare forms a package that we'll call *Animal Welfarism*.

Animal Welfarism is suggested by the views that we considered in Sect. 2 to motivate the Hybrid View. For example, the widely held idea that animal zoos are not morally bad as long as the animals are happy and satisfied, irrespective of any other considerations (e.g., their dignity or their freedom), suggests a view according to which animals' confinement is not intrinsically bad if it does not detract from their welfare (welfarism about animal value) and does not make the animals worse off if it does not negatively affect their mental states or frustrate their desires (subjectivism about animal welfare). Broadly, we believe, people who hold versions

Footnote 7 continued

Roger Crisp's distinction between enumerative and explanatory theories of welfare opens up the possibility that an individual's welfare consists, not in pleasure or desire-satisfaction, but rather in the things which the individual takes pleasure in for their own sake or in the things they desire non-instrumentally (Crisp 2006: 102–103). According to the traditional definitions, such views aren't versions of hedonism or preferentism, but they still count as versions of subjectivism according to our definition.

⁸ The Hybrid View can allow that the beauty of a painting hanging in the Louvre is intrinsically good. It would seem, then, that the Hybrid View should also allow that the beauty of a gazelle bounding in the Serengeti can be intrinsically good. But if the beauty of the gazelle is an animal state, then this is incompatible with welfarism about animals, which we are including as part of the Hybrid View. Thus, any plausible version of the Hybrid View will imply that not all states of animals are what we are calling *animal states*. Any fleshed-out version of the Hybrid View will require a criterion by which to distinguish animal states from other states. But different versions of the Hybrid View will require different criteria. So we cannot build a single definition of 'animal state' into our general characterization of the Hybrid View here.

of the Hybrid View will endorse the two commitments that comprise Animal Welfarism.

Next we want to discuss a further claim that will be endorsed by many but not all devotees of the Hybrid View. The two alternatives to scalar utilitarianism that we've discussed—maximizing utilitarianism and satisficing utilitarianism—are forms of *deontic utilitarianism*, according to which agents have obligations (Lang 2013). Deontic utilitarianism can be subdivided into directed utilitarianism and undirected utilitarianism.

Directed utilitarians hold that agents have obligations *to* others and that when an agent violates an obligation to an individual, the agent wrongs that individual. Undirected utilitarians hold that agents have obligations but do not have obligations to others. Undirected utilitarians allow that agents can act wrongly but deny that agents can wrong anyone.

Of course, a majority of deontologists—call them directed deontologists—also maintain that we have obligations to others. What distinguishes directed utilitarians from directed deontologists is that directed utilitarians maintain that our only fundamental (i.e., non-derived) obligations to others are obligations that involve and are explained in terms of their welfare or (equivalently) their interests, such that wronging an individual essentially involves leaving that individual worse off (relative to some baseline). In view of this, we propose that all directed utilitarians will endorse the following claim:

Directed Principle of Utility: An agent S wrongs an individual B only if S ϕ s such that (i) ϕ -ing is not optimific, and (ii) B is consequently worse off than B would have been had S acted rightly.

By contrast, a directed deontologist can say that S can wrong B without leaving B worse off than B would have been had S acted rightly; for example, by violating B's bodily integrity, rights, or dignity, or by using B merely as a means, or by violating egalitarian principles of justice in her treatment of B. And similarly, a directed deontologist can say that S can wrong B even if S acts optimifically (e.g., in the transplant case mentioned above).

We can now explain what it is to be a directed utilitarian *about animals*. Let's say that versions of the Hybrid View that incorporate directed utilitarianism about animals are versions of the Directed Hybrid View. We suggest that all defenders of a Directed Hybrid View will agree to the following:

Directed Principle of Animal Utility: If B is an animal, an agent S wrongs B only if S ϕ s such that (i) ϕ -ing is not optimific, and (ii) B is consequently worse off than B would be if S had acted rightly.

Note that the Directed Principle of Animal Utility is consistent with satisficing variants of directed utilitarianism about animals, because this principle is consistent with the view that we can in some cases act sub-optimifically, and thereby leave animals worse off than they would be if we had acted optimifically, without thus wronging those animals.

We've suggested that all defenders of the Hybrid View will be committed to (1) the Moral Significance Criterion, (2) the Negative Criterion, and (3) Animal

Welfarism (i.e., welfarism about animal value and subjectivism about animal welfare). And we've suggested that defenders of the Directed Hybrid View will sign on for (4) the Directed Principle of Animal Utility.

We have stopped short of offering a definition of the Hybrid View. Because of the wide diversity of utilitarian views, we can't define the Hybrid View without unduly constraining it, and a definition is not necessary for our purposes. The four theses just mentioned provide an informative picture of the contours of the Hybrid View. And they provide a characterization of the Hybrid View sufficient for our next task: to explain how the Hybrid View interacts with certain plausible claims about the moral significance of use.

4 The moral significance of use

The Hybrid View appears to be in tension with intuitively plausible claims about the moral significance of human use of animals. Consider a concrete case. The American military has used dogs for various purposes (e.g., as bomb-sniffers, as sentries) in many wars. Calhoun (2015) explains that these dogs have typically been classified as mere 'equipment,' and that when they have outlived their military usefulness, their fate has often been uncertain at best. For example, Calhoun says that at the end of the Vietnam War nearly 4000 military dogs were either euthanized or given away. It seems plausible that those who used those dogs had an obligation to look after their welfare, and that this obligation could not be discharged by euthanizing them or giving them away.

Of course, those 4000 dogs were presumably not the only dogs in Vietnam who could have benefited from American military assistance. Yet it is plausible that members of the American military had moral reasons to be *especially* concerned with the dogs they had used. Imagine: Lieutenant Mitchell is preparing to go home after the war. Rex is a German Shepherd who served Mitchell in a dangerous capacity for months as a bomb-sniffer. There is room on Mitchell's homeward-bound plane for one dog. If Mitchell puts Rex on that flight, Rex will happily live out the rest of his days with Mitchell's family; otherwise, the best that Lieutenant Mitchell can do for Rex is simply to release Rex onto the streets of Saigon, where Rex's life will be difficult and likely short.

Suppose that Mitchell does not rescue Rex and instead rescues a street dog with whom Mitchell has had no prior interaction. Such a decision seems to be a wrongful betrayal of Rex. By contrast, if Mitchell rescues Rex rather than the street dog, this would be unfortunate for the street dog but wouldn't be a betrayal.

We expect that many people who are attracted to the Hybrid View and who hold the sort whose views about zoos, research, and agriculture we discussed in Sect. 2 will feel the intuitive pull of these claims about Mitchell and Rex. They will not be intuitively skeptical of the view that use matters morally. In fact, we expect that many ordinary people will not perceive any conflict between their utilitarianism about animals and the view that use has direct moral significance.

Yet it is unclear whether the Hybrid View's utilitarianism about animals is truly consistent with the view that use has direct moral significance. After all, the mere

fact that an animal has been used in the past has no bearing on the welfare consequences for that animal of actions taken in the present (given Animal Welfarism) and those are the only consequences that utilitarians generally take to be morally relevant.⁹ So it is unclear whether the Hybrid View can allow that Mitchel has any special obligations to Rex.

However, in what follows, we'll show that utilitarianism about animals and the idea that use of animals matters morally can be brought into a surprising degree of harmony.

5 The Partiality Principle

We distinguish two principles that could govern the moral significance of use, beginning with the Partiality Principle. We'll show that, although this principle is not fully consistent with the Hybrid View, it is consistent with some key commitments of the Hybrid View.

Let's stipulate that a relationship between two individuals S_1 and S_2 is deontically significant when it affects the existence, content, or strength of the obligations that S_1 has to S_2 . And let's stipulate that S_1 is a moral stranger to S_2 iff S_1 and S_2 have no deontically significant relationships to one another. Consider:

The Partiality Principle: When a human S_1 uses an animal S_2 , S_1 is obligated to S_2 to act in a way that gives greater weight to S_2 's interests than to the equal interests of animals who are moral strangers to S_1 .

Many philosophers think our relationships with others can make partiality obligatory. Such views are usually framed in terms of human-to-human relationships, but they could be extended to cover human-to-animal relationships as well, in the manner suggested by the above Partiality Principle.¹⁰

The Partiality Principle requires a weighting function whereby an agent's use of an animal enhances the weight due to that animal's interests. Here's an example. First, an animal A's interest in an outcome O is assigned a base weight as a function of either the animal's desires regarding O, the hedonic stakes for the animal

⁹ A split-level or two-level utilitarian could argue that cultivating relationships with the animals we use promotes welfare in general and so we should, from a utilitarian perspective, generally approve of people who are unwilling to abandon animals they have used. Such a view fails to capture what, intuitively, would be morally problematic if Mitchell abandoned Rex. The issue is about Rex and Mitchell's relationship to him; it is not about whether Mitchell has internalized an optimistic set of motivations.

¹⁰ Keller (2013) provides a valuable taxonomy of partiality views: the projects view (defended by Bernard Williams among others), the relationships view (a default position in much recent partiality literature), and the individuals view (Keller's view). The Partiality Principle above is neutral between these three approaches. According to the Partiality Principle, given Mitchell's use relationship with Rex, Mitchell is obligated to show partiality toward Rex. This could be because of the role Rex plays in Mitchell's personal projects (as the projects view would have it), or because Mitchell's use relationship with Rex provides Mitchell with reasons to treat Rex in a special way (as the relationships view would have it), or because Mitchell's relationship with Rex enables Rex's intrinsic value to provide Mitchell with reasons to treat Rex in a special way (as the individuals view would have it). This illustrates that the Partiality Principle can be incorporated into a wide swathe of extant partiality views.

regarding O, or both. Second, A's interest in O is assigned a modified weight indexed to a given human agent H as a function of the base weight of A's interest in O together with facts about use of A by H.

This function—call it a *use function*—can take any number of different forms. In one option, use of A by H implies multiplying the base weight of A's interest by some factor X (where $X > 1$). X may vary depending on the sort of use at issue. It may be that if a dog is used for years in a hazardous role as bomb-sniffer, then X should be large, whereas if a dog has only been used as a sentry on one afternoon, then X should be small.

The Partiality Principle is incompatible with scalar utilitarianism about animals, as it implies that we have obligations. It is also incompatible with any undirected version of utilitarianism about animals, as it implies that humans can have obligations to animals.

More generally, the Partiality Principle is incompatible with any version of the Hybrid View, directed or undirected, because the Partiality Principle is inconsistent with the Moral Significance Criterion. In a case where failing to rescue Rex would not affect Mitchell or any other humans in any morally significant way, the Moral Significance Criterion implies that failing to rescue Rex is permissible if optimific—whereas the Partiality Principle says otherwise. According to the Partiality Principle, even if the street dog has a slightly greater interest in being rescued than Rex, Mitchell may still have an obligation to Rex to rescue Rex rather than the street dog (depending on the shape of the use function). Similarly, the Partiality Principle is inconsistent with the Directed Principle of Animal Utility, because the Partiality Principle implies that Mitchell can wrong Rex by leaving Rex behind even in cases where leaving Rex behind is optimific. It can likewise be seen that the Partiality Principle is inconsistent with the Negative Criterion.

However, the Partiality Principle is consistent with some of the core commitments of the Directed Hybrid View. First, it is consistent with Animal Welfarism. The views that only welfare is intrinsically good for animals and that animal welfare is subjective are fully consistent with the Partiality Principle's implication that we have an obligation to assign outsized weight to the interests of animals whom we have used.

And there is another, somewhat more subtle, consonance between the Partiality Principle and the Directed Hybrid View. To see it, first observe that the Directed Principle of Animal Utility implies the following principle:

Wronging Entails Making Worse Off: If B is an animal, an agent S wrongs B only if S ϕ s such that B is consequently worse off than B would be if S had acted rightly.

The Partiality Principle is fully consistent with Wronging Entails Making Worse Off. And Wronging Entails Making Worse Off can explain some of the attitudes that motivate the Hybrid View.

For example, consider the well-run zoo that we discussed in Sect. 2. Assume the animals are as happy and satisfied as they can be. If we also accept Animal Welfarism, then Wronging Entails Making Worse Off tells us that the animals in such a zoo cannot have been wronged by being confined and used in the zoo—no

matter how unfree, undignified, or unnatural their situation may be. Those who are inclined toward utilitarianism about animals will find this implication plausible.

And the Partiality Principle is consistent with that implication. The Partiality Principle says only that, given that the animals in a zoo are being used, zoo officials have an obligation to assign outsized weight to those animals' interests; and this is consistent with the idea that those animals cannot be wronged as long as their interests are optimally satisfied.

So while the Partiality Principle is inconsistent with the Hybrid View, it is consistent with motivating commitments of the Hybrid View, namely Animal Welfarism and Wronging Entails Making Worse Off.

That said, the Partiality Principle is inconsistent with another implication of the Hybrid View. The Hybrid View implies that, holding constant effects on humans, it is always permissible to be *impartial* to animals: that is, it is not wrong to give equal weight to the interests of all animals who are affected by our actions. Let us call this idea *Animal Impartiality*.

We may distinguish two forms of animal impartiality. According to *Strong Animal Impartiality*, we are morally required to act in a way that gives equal moral weight to the equal interests of animals. According to *Weak Animal Impartiality*, we are *not* morally required *not* to act in a way that gives equal moral weight to the equal interests of animals. All utilitarians about animals will endorse Weak Animal Impartiality; maximizing utilitarians about animals will endorse Strong Animal Impartiality.

Both Strong and Weak Animal Impartiality are inconsistent with the Partiality Principle, given that the Partiality Principle obligates us to assign greater weight to the interests of animals with whom we have a use relationship. To the extent, then, that one is attracted to Animal Impartiality in either its weak or strong versions, the Partiality Principle will seem unsatisfactory.

It is worth noting, however, that an agent-relative notion of impartiality is consistent with the Partiality Principle: if two people stand in the same relationship to the animals that they each use, then one person's relationship to the animal they use is neither more nor less morally significant than the other person's. Just as the Smith Lab ought to give outsized weight to the interests of the animals they use, so too should the Jones Lab give outsized weight to the interests of the animals *they* use. Although such agent-relative impartiality falls short of the kind of impartiality that utilitarians endorse, it reduces the gap between the Partiality Principle and the Hybrid View.

6 The Strengthening Principle

The Partiality Principle is not the only way of understanding the moral significance of animal use. Another way is to think of use as strengthening obligations, rather than as creating new obligations or affecting the content of one's obligations. For instance, a Hybrid View that incorporates maximizing utilitarianism about animals might endorse the following principle:

The Strengthening Principle: When a human S_1 uses an animal S_2 , this strengthens S_1 's obligation to S_2 to act in a way that gives equal weight to S_2 's interests as to the equal interests of other animals (e.g., animals who are moral strangers to S_1).

The Strengthening Principle is consistent with both Strong and Weak Animal Impartiality: it does not say that we are obligated to give greater weight to the interests of animals we have used than to the equal interests of animals who are moral strangers. All it says is that it is more seriously wrong for a human user to fail to treat an animal usee as if her interests are of equal weight as the equal interests of other animals, than it is to fail to treat an animal who is a moral stranger to S_1 in such a way.

To see the idea of the Strengthening Principle, consider two cases:

Bad Herd Management: Farmer Jeb has two groups of sheep. The first has been vaccinated to prevent a deadly and painful disease. It's time to vaccinate the second group, but because Farmer Jeb is careless, he vaccinates the first group again and the second group not at all. Consequently, several sheep in the second group become sick and painfully die.

Bad Charity Decision: Lisa has money to donate to one of two animal charities. The first would do significantly more good for animals than the second. Lisa could easily find this out, but she does not bother. The second charity has a snappier-sounding name, so she sends the money to the second charity, with the result that several sheep, who would have been fine if she had donated to the first charity, become sick and painfully die.

Let us stipulate that the consequences for animal welfare are the same in each case. Deontic versions of the Hybrid View will imply that both Farmer Jeb and Lisa act wrongly: they fail to maximize utility and fail to meet any plausible threshold for doing at least good enough. The Strengthening Principle is consistent with this judgment. But the Strengthening Principle further implies that Jeb acts *more seriously* wrongly than Lisa (*ceteris paribus*¹¹) because Jeb has a use relationship with his sheep whereas Lisa has no such relationship with the animals who stand to benefit from her charitable giving. That's intuitively plausible: while Lisa has failed to do the most good for animals she can do with her charitable giving, what she has done does not seem as seriously wrong as what Jeb has done.

What are the implications of the fact that one obligation is stronger than another? At least two. First, agents deserves harsher blame for violating stronger obligations (*ceteris paribus*). Second, we should prioritize stronger obligations over weaker ones, such that if an akratic agent foresees that she will not do all that she is obligated to do, then she had better shirk her weaker obligations rather than her stronger ones.

¹¹ The Strengthening Principle only says that use strengthens obligations; it does not say that use is the only factor relevant to the strength of an obligation. Thus it is compatible with the Strengthening Principle that an obligation to a stranger can be stronger than an obligation to a usee. We thank an anonymous referee for pushing us to make this point.

To illustrate, suppose you are drunk at a party. You can get a taxi ride home, ask your host to let you sleep over, or drunkenly drive home. Getting a taxi ride would be best, but you are stingy and will not pay for it (and no one else will either). Suppose asking to sleep over would violate a comparatively weak obligation not to abuse your host's hospitality and drunkenly driving home would violate a comparatively strong obligation not to endanger innocent lives. Then you'd better ask to sleep over rather than drunkenly drive home.

So, the Strengthening Principle has two important implications. First, when we use an animal and yet fail to act in a way that gives equal weight to that animal's interests as to the equal interests of other animals, then we deserve especially harsh blame. And second, if we will fail to act in a way that gives equal weight to the equal interests of all animals affected by our actions, then we should prioritize giving equal weight to the interests of animals we use.

Like the Partiality Principle, the Strengthening Principle is inconsistent with scalar and undirected versions of the Hybrid View. But the Strengthening Principle is consistent with *all* of the commitments of the Directed Hybrid View that we have identified—the Moral Significance Criterion, the Negative Criterion, Animal Welfarism, the Directed Principle of Animal Utility, Weak Animal Impartiality, and, for maximizing views, Strong Animal Impartiality.

And yet the Strengthening Principle affirms that the use of animals has direct moral significance. Given the Strengthening Principle, the use of animals matters morally because use strengthens our obligations to animals, although use does not change the content of those obligations. Thus the availability of the Strengthening Principle establishes the striking result that we can be full-fledged devotees of the Hybrid View while also capturing the commonsense idea that our obligations to animals we have used differ in an important way from our obligations to animals who are moral strangers to us.

That said, the Strengthening Principle is incompatible with one more commitment that we might expect to be widely (though not universally) shared by those who hold a Hybrid View. Many philosophers, utilitarians and deontologists alike, take the distinction between utilitarian and non-utilitarian ethical theories to ultimately rest on disagreement about the relationship between the deontic and the evaluative, or as it is commonly phrased, between “the right and the good” (Rawls 1971: 24–25). In this vein, utilitarianism is usually understood to be an expression of the Moorean idea that the deontic supervenes on the evaluative, i.e., there can be no difference in the obligations faced by agents in two cases without a difference in the value of the consequences of the actions available to the agents. Now consider a restricted variant that applies only to cases that only involve animals:

Restricted Supervenience: For any pair of cases C_1 and C_2 in which

- (i) an agent S must decide whether to ϕ ,
- (ii) S 's ϕ -ing affects animals but not humans in morally significant ways, and
- (iii) there is no difference in the intrinsic value of the consequences of the actions available to S in C_1 and C_2 ,

there can be no difference in S 's obligations in C_1 and C_2 .

Given that Restricted Supervenience is an animal-specific version of a thesis that is widely accepted by utilitarians, we may expect it to be attractive to many who endorse the Hybrid View. Yet Restricted Supervenience is not consistent with the Strengthening Principle: restricted Supervenience precludes the possibility that the strength of our obligations to animals can be affected by use relationships in the way required by the Strengthening Principle. (It should be noted that the Partiality Principle, which implies the existence of a new obligation, is also incompatible with Restricted Supervenience.) Those who are attracted to the Hybrid View and to the moral significance of animal use may do best to opt for the Strengthening Principle, as the Strengthening Principle at least allows that the content of our obligations to animals supervenes on the evaluative, even if the strength of those obligations does not. But it may still require a modification to their view about the metaphysical relationship between the evaluative and the deontic.

7 Motivating the Strengthening Principle

Versions of the Hybrid View that include the Strengthening Principle have an unusual implication: the *strength* of our obligations to animals is affected by our use relationships with them, but the *content* of those obligations is not. Rather, the content is determined exclusively by facts about animals' interests.

This may seem implausible. Worse, it may seem ad hoc: it may seem that we are proposing the Strengthening Principle for no reason other than as a logical exercise in reconciling the moral significance of use with the Hybrid View.

However, there are cases where it is plausible that our relationships with others affects the strength but not the content of our obligations. Suppose that you are a battlefield surgeon. Ten wounded soldiers have just been delivered to your station. To your horror, you discover that your own daughter is among the wounded. The other nine are strangers.

Your obligation as a battlefield surgeon is to apportion your medical resources and time in such a way as to save as many soldiers' lives as possible. To do this job, you are obligated to treat your daughter like all the others: you must act as if she is a stranger. If one soldier has an excellent chance of survival if that soldier receives medical attention and will die if left unattended, while another has almost no chance of survival even with your attention, then you must attend to the first before attending to the second—even if the second is your daughter.

In this situation it is plausible that (A) your obligation as a battlefield surgeon is to treat all of the soldiers' interests equally, i.e., you must not exhibit partiality toward any one of them. Nevertheless, it is also plausible that (B) if you were to exhibit partiality toward one of the strangers at the expense of your daughter (e.g., by attending to one of the strangers rather than your daughter even though your daughter is more likely to benefit from medical attention than the stranger) then this would be more seriously wrong than if you were to exhibit the same sort of partiality toward your daughter at the expense of one of the strangers.

Although there may appear to be tension between (A) and (B), there is also a straightforward way to reconcile (A) with (B): the content of your obligations in this

triage case is determined by a different function than the strength of those obligations. Specifically, the content of your obligations is a function of the base weight (see Sect. 5) of all of the soldiers' interests, but the strength of your obligations is a function of the base weight of soldiers' interests adjusted by your parent-child relationship with one of those soldiers.

In this case, the idea that the content of your obligations is determined by a different function than the strength of your obligations is not merely an ad hoc exercise in reconciling seemingly opposed theoretical commitments. Rather, that idea is quite directly motivated by (A) and (B).

The relevance of this for the Strengthening Principle emerges as follows. According to utilitarianism, it is not just in triage cases that the content of our obligations is a function of the interests of those who are affected by our actions. Thoroughgoing utilitarians hold that *all* of our obligations are like that. For thoroughgoing utilitarians, the whole world is essentially a vast triage case.

Of course, defenders of the Hybrid View are not thoroughgoing utilitarians. They are merely utilitarians about animals. Utilitarianism about animals implies that the content of our obligations to animals is a function of the base weight of animals' interests. So we can speculate that people who are drawn to utilitarianism about animals (Sect. 2) will find it plausible that the content of our obligations to animals is a function of the base weight of animals' interests. But many of these people may also find it plausible that the strength of our obligations to animals is affected by our use relationships to animals, such that it is more seriously wrong to discount the interests of animals one is using or has used than to discount the interests of animals one has no relationships with.

And so it seems plausible that the main requirement of the Strengthening Principle—that our relationships with animals affect the strength of our obligations to animals without affecting the content of our obligations to animals—can be independently motivated and is therefore not ad hoc, just as the reconciliation of (A) with (B) above is not ad hoc.

However, we must admit that the Strengthening Principle may not successfully capture all of our intuitions about the moral significance of use. Consider two versions of the Rex and Mitchell case. In the first, a street dog has a slightly greater interest in being rescued than Rex does. In the second, Rex has a slightly greater interest in being rescued than the street dog does.

Intuitively, in virtue of Mitchell's use relationship with Rex, it would be more seriously wrong to show partiality toward the street dog at Rex's expense in the second case than to show partiality toward Rex at the street dog's expense in the first case. It is an advantage of the Strengthening Principle that it can explain and support this intuition.

But some may have the further intuition that it would *not be wrong at all* for Mitchell to show partiality toward Rex at the street dog's expense in the first case, and that it would be wrong not to do so—even though the optimific course of action is (ex hypothesi) to rescue the street dog rather than Rex. The Strengthening Principle cannot support this further intuition. Those who are subject to this further intuition ought to prefer the Partiality Principle over the Strengthening Principle.

But for those who are not subject to this further intuition, the Strengthening Principle may be seen to have significant intuitive advantages.

8 Conclusion

Many people appear to be utilitarians about animals but find utilitarianism about humans objectionable (Sect. 2). We've identified various commitments of utilitarianism about animals: the Moral Significance Criterion, the Negative Criterion, Animal Welfarism, the Directed Principle of Animal Utility, Wronging Entails Making Worse Off, Strong and Weak Animal Impartiality, and Restricted Supervenience (Sects. 3, 5, 6). We hope this fine-grained approach to the question of what it is to be a utilitarian about animals will provide a taxonomy of views that will foster future debate about the implications of utilitarianism about animals and about whether utilitarianism about animals is, in fact, true.

The Hybrid View faces a challenge in accommodating the moral significance of animal use (Sect. 4). It's intuitively plausible that using an animal is a deontically significant relationship that affects our obligations to the animals so used. To explore whether this idea is consistent with utilitarianism about animals and with other theses that utilitarians find attractive, we've distinguished between two ways of understanding the moral significance of use: the Partiality Principle and the Strengthening Principle.

The Partiality Principle is compatible with Animal Welfarism and with Wronging Entails Making Worse Off, but is inconsistent with several other main commitments of the Hybrid View (Sect. 5). By contrast, the Strengthening Principle can be incorporated into a directed version of the Hybrid View (Sect. 6). Thus, in addition to being independently motivated, the Strengthening Principle represents a way for those who want to be utilitarians about animals to nonetheless acknowledge that our use of animals has direct moral significance.

This is a striking conclusion in light of traditional utilitarian resistance to the idea that special relationships can be of direct moral significance. And importantly, our reconciliation is compatible with Animal Welfarism and thus does not require endorsing a pluralistic, non-subjectivist axiology, unlike Peter Railton's classic attempt to accommodate special relationships within a consequentialist framework (1984: 148–150). Such an axiology may be plausible for humans, but many will find it implausible for animals.

Many people find traditional utilitarian approaches to animals unsatisfactory because they seem to fail to do justice to the moral importance of our personal relationships and interactions with particular animals. Some such people, though, may still be reluctant to endorse the idea that animals have rights. Such people might find the Directed Hybrid View combined with the Strengthening Principle attractive as an intermediate position.

Our aim in this paper has been to characterize the Hybrid View and explore its commitments and its compatibility with intuitive claims about use. We cannot offer a complete evaluation of the Hybrid View here—that is a project that goes well beyond the scope of this paper. But before we finish, we want to mention some

issues that will have to be addressed in future discussions about the plausibility of the Hybrid View.

First, it may be argued that the Hybrid View entails a moral hierarchy in which animals are relegated to a lower status than humans. Utilitarianism about animals allows and (in some versions) requires the interests of animals to be traded off in ways that would be prohibited by deontological views about humans. In other words, the Hybrid View implies that humans are *morally inviolable* to a greater degree than animals (Kamm 2007: 26–30). And it may be argued that such a view objectionably elevates humans above animals.

However, it is not obvious that assigning a greater degree of moral inviolability to a class of individuals means elevating that class above others. Moral inviolability is not always desirable. In fact, it can be burdensome or even fatal to be a member of a class of beings who are regarded as morally inviolable: for example, a deontological prohibition against killing one to save many can be disastrous for you if you are among the many (Kagan 1991: 919–920). Given this, the question whether the Hybrid View assigns animals to a lower status, or a higher status, or simply to a different status requires further investigation.

Second, a great many have argued that there is no morally important feature that all humans have and that all animals lack. If they are right, then the Hybrid View relies on a morally arbitrary distinction and so is implausible. Defenders of the Hybrid View will need to either argue that the human/animal distinction is not morally arbitrary, or modify the view so that it relies on a different distinction (e.g., a distinction between persons and non-persons). These and many other issues will need to be addressed in order to adequately evaluate the Hybrid View.

Acknowledgements We're indebted to Stephanie Collins, Roger Crisp, Bob Fischer, Tyler John, Wojtek Kaftanski, Josh Kissel, Ole Koksvik, the Koronczyks, Emilian Mihailov, Paul Mikhail-Podosky, Josh Mund, Tyler Paytas, Richard Rowland, Russ Shafer-Landau, and an anonymous reviewer for this journal.

References

- Bentham, J. (1789). An introduction to the principles of morals and legislation. In J. Burns & H. Hart (Eds.), London: Methuen, 1970.
- Brandt, R. B. (1988). Fairness to indirect optimific theories in ethics. *Ethics*, 98, 341–360.
- Calhoun, C. (2015). But what about the animals? In Timmons & Johnson (Eds.), *Reason, value, and respect: Kantian themes from the philosophy of Thomas E. Hill, Jr.* (pp. 194–211). Oxford: Oxford University Press.
- Caviola, L., Everett, J., Teperman, E., Savulescu, J., & Faber, N. (in preparation). Utilitarianism for animals, kantianism for people: Speciesism in sacrificial moral dilemmas.
- Cohen, C. (1986). The case for the use of animals in biomedical research. *The New England Journal of Medicine*, 315(14), 785–870.
- Cohen, C., & Regan, T. (2001). *The animal rights debate*. Boston: Rowman and Littlefield Publishers Inc.
- Crisp, R. (2006). *Reasons and the good*. Oxford: Clarendon Press.
- Crisp, R. (m.s.). Partiality and the dualism of practical reason, unpublished manuscript.
- De Lazari-Radek, K., & Singer, P. (2014). *The point of view of the universe*. Oxford: Oxford University Press.
- Gennet, E., & Altavilla, A. (2016). Paediatric research under the new EU Regulation on clinical trials: Old issues new challenges. *European Journal of Health Law*, 23(4), 325–349.
- Gruen, L. (2011). *Ethics and animals: An introduction*. Cambridge: Cambridge University Press.
- Hare, R. M. (1981). *Moral thinking: Its levels, method, and point*. Oxford: Clarendon University Press.

- Institute of Medicine. (2011). *Chimpanzees in biomedical and behavioral research: Assessing the necessity*. Washington, D.C.: National Academies Press.
- Jamieson, D., & Elliot, R. (2009). Progressive consequentialism. *Philosophical Perspectives*, 23(1), 241–251.
- Kagan, S. (1989). *The limits of morality*. Oxford: Clarendon University Press.
- Kagan, S. (1991). Replies to my critics. *Philosophy and Phenomenological Research*, 51(4), 919–928.
- Kagan, S. (1998). *Normative Ethics*. Boulder, CO: Westview Press.
- Kamm, F. (2001). *Mortality, morality volume II: Rights, obligations, and status*. Oxford: Oxford University Press.
- Kamm, F. (2007). *Intricate ethics: Rights, responsibilities, and permissible harm*. Oxford: Oxford University Press.
- Keller, S. (2013). *Partiality*. Princeton: Princeton University Press.
- Knight, A. (2011). *The costs and benefits of animal experiments*. London: Palgrave Macmillan.
- Lang, G. (2013). Should utilitarianism be scalar? *Utilitas*, 25, 80–95.
- McMahan, J. (2002). *The ethics of killing: Problems at the margins of life*. Oxford: Oxford University Press.
- McMahan, J. (2008). Eating animals the nice way. *Daedalus*, 137, 66–76.
- McMahan, J. (2016). The comparative badness for animals of suffering and death. In T. Visak & R. Gardner (Eds.), *The ethics of killing animals*. Oxford: Oxford University Press.
- Mill, J. S. (1861). *Utilitarianism*. Indianapolis: Hackett Publishing Company.
- Mogensen, A. (2016). Should we prevent optimific wrongs? *Utilitas*, 28, 215–226.
- Moore, G. E. (1903). *Principia ethica*. Cambridge: Cambridge University Press.
- Norcross, A. (2006). The scalar approach to utilitarianism. In H. R. West (Ed.), *The Blackwell Guide to Mill's utilitarianism*. Malden, MA: Blackwell.
- Nozick, R. (1974). *Anarchy, state, and Utopia*. New York: Basic Books.
- Nuffield Council on Bioethics. (2005). *The ethics of research involving animals*. London.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Clarendon Press.
- Railton, P. (1984). Alienation, consequentialism, and the demands of morality. *Philosophy & Public Affairs*, 13(2), 134–171.
- Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Belknap Press.
- Regan, T. (1983). *The case for animal rights*. Berkeley: University of California Press.
- Ross, W. D. (1930). *The right and the good*. Oxford: Oxford University Press.
- Ross, L. (2005). Lessons to be learned from the 407 process. *Health Matrix*, 15, 401–421.
- Sandøe, P. (2003). Farm animal welfare: The interaction of ethical questions and animal welfare science. *Animal Welfare*, 12, 469–478.
- Sidgwick, H. (1907). *The methods of ethics* (7th ed.). Indianapolis: Hackett Publishing Company.
- Sinclair, T. (2017). Preventing optimific wrongings. *Utilitas*, 29, 453–473.
- Singer, P. (2011). *Practical ethics* (3rd ed.). Cambridge: Cambridge University Press.
- Sinhbabu, N. (2018). Scalar consequentialism the right way. *Philosophical Studies*, 175, 3131–3144.
- Slote, M., & Pettit, P. (1984). Satisficing consequentialism. *Proceedings of the Aristotelian Society*, 58, 139–176.
- Smart, J. J. C., & Williams, B. (1973). *For and against*. Utilitarianism: Cambridge University Press.
- Thomson, J. J. (1994). Goodness and utilitarianism. *Proceedings and Addresses of the American Philosophical Association*, 67(4), 157–172.
- Varner, G. (1994). The prospects for consensus and convergence in the animal rights debate. *Hastings Center Report*, 24(1), 24–28.
- Varner, G. (2012). *Personhood, ethics, and animal cognition: Situating animals in hare's two-level utilitarianism*. Oxford: Oxford University Press.
- Woodard, C. (2013). Classifying theories of welfare. *Philosophical Studies*, 165(3), 787–803.