Utilitas (2020), **32**, 276–293
doi:10.1017/S0953820819000542

CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

Valuing humane lives in two-level utilitarianism

Nicolas Delon*

New College of Florida
*Corresponding author. E-mail: nd.delon@gmail.com

(First published online 20 January 2020)

Abstract

I examine the two-level utilitarian case for humane animal agriculture (by R. M. Hare and Gary Varner) and argue that it fails on its own terms. The case states that, at the 'intuitive level' of moral thinking, we can justify raising and killing animals for food, regarding them as replaceable, while treating them with respect. I show that two-level utilitarianism supports, instead, alternatives to animal agriculture. First, the case for humane animal agriculture does not follow from a commitment to two-level utilitarianism combined with a commitment to respecting animal lives. Second, the two-level utilitarian case falls prey to a compartmentalization problem and cannot uphold both respect and replaceability. What I call 'humane lives' are not appropriately valued by the lights of two-level utilitarianism itself.

1. Introduction

Two-level utilitarianism (TLU) is a form of indirect consequentialism. Its structure is simple. It consists of a criterion of rightness (act-utilitarianism) and a decision procedure (rules or principles that agents can employ in everyday circumstances), which are respectively dubbed the "critical" and "intuitive" levels of moral thinking (Hare 1981). Can TLU satisfy those who reject standard consequentialist thinking about animals (say, because it regards them as replaceable)? My goal in this article is to revive an objection levelled by Bernard Williams (1985; 1988) against Hare's theory[1] and to bring it to bear on Gary Varner's (2012) recent Harean defense of humane agriculture. My central concern is whether the lives of humanely raised and killed animals, *humane lives*, are appropriately valued in TLU. I press three lines of argument. First, I show that Hare and Varner, through their commitment to replaceability, weaken important non-utilitarian aspects they think are justifiably embedded in intuitive thinking. Then I argue that Varner's theory fails, by its own lights, to support humane agriculture. Finally, I explicate the internal tension of TLU, the problem of compartmentalization, and argue that it generates incoherent psychological attitudes that compound the problem: it cannot uphold respect-cum-replaceability. The article proceeds as follows. In section 2, I reconstruct Varner's application of Hare's theory to animal ethics. Section 3 unpacks some assumptions of TLU's case for humane slaughter-based agriculture, introducing the replaceability argument. Section 4 describes Varner's

[1]For discussion of the Hare–Williams debate, see Levy (1994) and Miller (2014).

© Cambridge University Press 2020

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

framework of "visions of humane sustainability" and argues that we can derive commitments against animal agriculture from the same framework. Section 5 brings the problem of compartmentalization, and Williams' objection, to bear on the debate, in light of recent psychological research.

2. Two-level utilitarianism and animals

2.1. Rules and dual-process theory

Hare is a total act-utilitarian,[2] but typical human agents are more likely to maximize aggregate utility (the well-being of all actual and possible beings), not by regularly relying on explicit consequentialist thinking, but by employing some hybrid decision procedure. According to TLU, that procedure consists in "prima facie principles" (Hare 1981: 38), or rules, that we should internalize such that violating them, or considering doing so, causes in us feelings of guilt and compunction (ibid. 29–31). In other words, when we have internalized the correct rules we are disposed not only to act in certain non-utilitarian ways but also to experience feelings that express our commitment to such rules.

Following Varner, call the rules of intuitive-level system *ILS rules*. The substantive thesis that we should often rely on simple moral heuristics follows from the empirical thesis that we are cognitively limited. So, TLU can both *explain* our intuitive reactions to standard utilitarian cases and *justify* some of those reactions in ordinary contexts – on the ground that having such reactions tends to promote utility. Hare provides two stylized character descriptions of the two extremes between which moral thinking takes place. The "archangel" is not subject to normal cognitive limitations; they can think exclusively critically. They would know, at any given time, which of the options available to them would maximize utility. In contrast, the "prole" has cognitive limitations that are severe enough that they would do best to think exclusively intuitively (Hare 1981: 44–46). Most of us are neither archangels nor proles; "we all share the characteristics of both to limited and varying degrees and at different times" (ibid. 45). Most of our everyday decisions should be made at the intuitive level, but some unfamiliar contexts will call for more deliberation. Critical thinking is deployed to reflect on the rules themselves "in a cool hour" – on whether and how they should be overridden, revised or replaced. For the rules that we internalize can conflict or fail to provide sufficient guidance or need to be updated, especially in novel cases. So, when we have enough information, time and cognitive resources, critical thinking supplements rules to determine the best course of action. For instance, we might be justified in violating the rights of some people to save a large number of lives. Because such circumstances are exceptional, good ILS rules include strong constraints against harming persons. Rules that have a "deontological flavor" can be more likely to do good given what we know about the world and human psychology.

Dual-process theories posit two distinct cognitive systems and/or types of processes. "System 1" (fast, automatic, emotional, stereotypic, subconscious) and "system 2" (slow, effortful, deliberative, rational, conscious) (Kahneman 2011) roughly correspond to Hare's intuitive and critical levels. Joshua Greene draws a helpful analogy: the two systems are like the "point-and-shoot, automatic" settings and "manual mode" of a digital camera (2013: 132–43). One is used frequently, the other less so. One is more

[2]The structure of TLU is compatible with other forms of utilitarianism (average, person-affecting, hedonistic, etc.). Varner seems to endorse a hybrid of desire-fulfillment theory and hedonism about well-being. But this is orthogonal to my critique.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

emotional, the other more cognitive and deliberative. System 1 is more efficient and inflexible, and the "central tension" of our dual-process brains is to negotiate trade-offs between efficiency and flexibility. Manual mode thinking, which allows for controlled, conscious reasoning, and which we all share, is predisposed to utilitarian thinking (ibid. 198). And it's meant to solve problems that automatic settings can't solve reliably – problems such as global poverty and climate change that are unfamiliar to the brains our ancestors evolved.

The two levels are thus two sides of a unified psychology that sometimes conflict but can complement each other. In fact, the descriptive part of Hare's theory *predicts* that our intuitions will resist the verdicts of the critical level, and in most circumstances, this is as it should be. At the same time, dual-process theorists tell us that our intuitions will often misfire, in which case conscious, controlled reasoning should step in. On Varner's interpretation of TLU, ILS rules include rules of personal morality, common morality, professional ethics, and the law. ILS rules work best most of the time because they yield quick and easy answers. General though they are, they can be fine-grained enough to accommodate different individual "temperaments" and "capacities" (which fall under personal morality) as well as the particular circumstances of different societies (e.g., through law and common morality). For different individuals or societies different sets of rules apply. Rules should be general enough to release individuals of the burden of critical thinking, yet specific enough to accommodate particular circumstances.

Varner believes that TLU is uniquely equipped to avert both standardly consequentialist welfarism and anti-utilitarian abolitionism about animal use, by articulating a decision procedure involving non-utilitarian elements and critical utilitarian reasoning. Part of that decision procedure includes different rules for different kinds of animals.

2.2. Categories and moral status

Varner divides animals into *persons*, *near-persons*, and *merely sentient*, and outlines the possible sets of rules that should govern our interactions with them. The three distinct, if fuzzy, tiers reflect a spectrum of cognitive complexity and yield a corresponding three-tier account of moral status.

Merely sentient

Most if not all vertebrates, as well as cephalopods, are most likely to be sentient, i.e., have "the capacity for phenomenally conscious suffering and/or enjoyment" (Varner 2012: 108). Most sentient animals are "merely sentient" (i.e., lack the capacities instantiated by the other two tiers) and most animals currently raised and slaughtered for food fall under this category.

Persons

A person is a rational, self-conscious, autonomous being with a biographical sense of self. In order to construct self-narratives, one must have concepts (of self, birth, death, and personality), hence a natural language (chapter 6). Persons can tell life stories, desire to be a certain kind of person, and have a unique interest in how their "lives-as-a-whole" go. This implies understanding oneself as having desires and being the subject of a story pursuing them, but also having desires about one's desires. These capacities give the lives of persons greater moral significance, and their higher

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

moral status is translated into rights or side-constraints embedded in ILS rules. According to Varner, we currently have no good evidence that nonhuman species qualify for personhood.

Near-persons

Some animals sit in between the merely sentient and persons. They have autonoetic awareness, or a sense of self over time (chapter 7). Great apes, dolphins, elephants, scrub jays and other corvids are near-persons; rats and monkeys might be. While near-persons cannot tell stories about themselves, they can engage in forward- and backward-looking thinking. Some can even perform "mental time travel," that is, anticipate and remember future and past states of the world and their own experiences in context (i.e., episodic memory; Tulving 1985). Some can recognize themselves in mirrors and/or engage in some form of mindreading. Autonoetic awareness allows individuals to re-experience past events and to make future plans that matter to them. Their lives are thus more "morally charged," containing more opportunities for well-being. These animals can therefore be harmed or benefitted in more ways.[3]

Varner's theory recommends incremental yet significant reforms of animal agriculture. Industrial animal agriculture and fishing are morally impermissible because they fail to respect sentient animals' most basic interests, and this failure is not justifiable by the human interests that are thereby satisfied. The principle "Don't kill sentient animals unnecessarily" should be incorporated into the ILS rules of all societies (Varner 2012: 229), along with a principle that sentient animals deserve respect (in a rather thin sense of direct consideration for their interests). But the principle does not specify exactly *when* killing animals is permissible or required. An ILS prohibition against unnecessary killing does not apply to practices that are on balance optimific. Indeed, an absolute ban on raising and killing farm animals at the intuitive level would be non-optimific, according to Hare and Varner. To see why, let us consider the replaceability argument, to which both are committed.

3. The replaceability argument

Some early version of the argument can be traced at least back to Bentham (1907) [1789] and Stephen (1896). These versions were later dubbed "the logic of the larder" by the early animal rights writer Henry Salt (1914), and consequentialists have frequently endorsed a version of the argument, as have many "conscientious omnivores" such as Michael Pollan (see Delon 2016). Many people who condemn inflicting unnecessary suffering upon animals believe that the painless killing of animals may be permissible. When slaughtered without pain, fear or distress, the thought goes, animals are not harmed by death in a morally significant way. For many utilitarians and others too, suffering and death are morally distinct; many animals lack a morally significant interest in continuing to live. Hare (1999) and Varner (also see Singer 2011; Singer 2016; de Lazari-Radek and Singer 2014) argue that existence can be a benefit if one's life is worth living, and a good life can offset a premature death, for animals who would not exist otherwise. It follows that breeding happy animals in order to

[3]Varner's recognition of near-persons has limited scope: 98% of the animals that American society interacts with are farm animals (Wolfson and Sullivan 2004). Around 70 billion land animals are killed for food globally every year, trillions if we count aquatic species.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

kill and eat them can be good on balance, for consumers and for animals. In a slogan: we should not just breed *happier* animals but also *more* happy animals, since the preferences of possible beings also count (Hare 1993; 1999). Of course, there is a limit to how many more happy lives we can create, but this doesn't block the replaceability argument from allowing, at least in principle, breeding and killing animals in certain circumstances – even Peter Singer agrees.[4] Hare (1999: 239) wrote that "happy existing people are certainly glad they exist, and so are presumably comparing their existence with a possible non-existence." It is better "for an animal to have a happy life, even if it is a short one, than no life at all." It followed, he argued, that "organic" husbandry is permissible,[5] perhaps even required on marginal lands where crops cannot be grown.

The replaceability argument has generated much debate (Delon 2016; Višak and Garner 2016). Not all consequentialists or utilitarians accept it (John and Sebo, forthcoming; Višak 2016). It makes controversial assumptions about the comparative value of existence relative to non-existence, the possibility of assigning zero welfare to non-existence, as well as the significance of the harm of death to animals that lack a sense of their future. Precisely because it is controversial, I suggest that ILS rules should not, as a general matter, presuppose it. Admittedly, the argument is deployed at the critical level. However, even if it were to remain "esoteric," and thus hidden from most people as a *rationale* for animal agriculture, there is no practical way of preventing people from regarding animals as replaceable if ILS rules make it permissible to treat them as substitutable food sources. The point is not that ILS rules should be based on consensus; rather, that what presuppositions we put into ILS rules affects how efficiently they can foster the right sort of attitudes to animals.

With TLU's case for humane agriculture, and a first objection, in the background, I build up a two-part critique in sections 4–5. I start with Varner's three-tier framework for assessing farming techniques in light of humane sustainability, and explain how the framework can be used to defend alternatives to animal agriculture. In section 5, I will address the compartmentalization problem that arises from the appeal to replaceability.

4. Utopian visions

In this section, I discuss Varner's framework for assessing farming techniques in context (2012: ch. 10). The framework is key to articulating the moral demands embedded at the intuitive level. My central claim is that the same framework can be used to

[4]Singer (2011) gives the benefit of the doubt to most of the animals we currently raise for food (more so than in 1993), crediting even birds and fish with some capacity for anticipation and future-directed preferences. In 2011, Singer rejected the replaceability of persons but accepted it, in principle, for nonpersons; his recent endorsement of hedonism (2016; de Lazari-Radek and Singer 2014) pushed him to endorse the argument in its wide scope.

[5]Is it also *required*? This depends on whether not eating animals at all would be unjustifiably suboptimific. Varner sees animal activists' actions as "supererogatory" (2012: 270), so it's at least not forbidden as a matter of *personal* morality. Still, it would be wrong to enforce it as a matter of *policy*. Moreover, both veganism and factory farming deny (possible and actual) animals happy lives, running afoul of replaceability. So, Varner, like Hare, seems committed to preserving at least some animal husbandry. Certainly, whether we should *eat* animals is not entailed by the reasons we have to *breed* them. We might "harvest" them humanely to prevent them from dying worse deaths – no doubt a lesser harm than killing them prematurely for human consumption. However, the replaceability argument assumes that it is not economically feasible to raise happy animals without killing them in their prime and selling their flesh or their products. Thanks to Dale Miller for pressing these points.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

promote very different visions of humane sustainability, and humane agriculture may not be the one a Harean should endorse. I will offer some empirical and some speculative hypotheses in support of the claim that TLU requires promoting the idea that we stop eating animals.

4.1. Visions of humane sustainability

Varner lays out three types of visions that each come with different standards of assessment.

Pre-lapsarian visions are the visions of "ancient societies that survived relatively unchanged for millennia," which Varner calls so "because environmentalists often venerate such societies and assume that their descriptively sustainable practices also met whatever norms are implicit in their conception of sustainability" (2012: 260). Native Americans were once justified in hunting bison in ways that caused great suffering and sometimes prevented them from using all of the animals killed. But modern, industrialized societies may no longer permissibly do so given the available technology and our understanding of animal welfare. Because ILS rules display a "benign form of relativism", "the drive hunts of pre-contact Native Americans may well have been an example of humane sustainability *for persons in their circumstances* … [but] surely modern humans would not be justified in using such inefficient and relatively inhumane hunting techniques" (ibid. 262–63). TLU explains why a given practice can be wrong in our circumstances but permissible in others.

Contemporary visions are attempts to improve animal welfare using contemporary science and engineering "constrained, at least generally, by contemporary consumer preferences and economic realities" (Varner 2012: 262). Animal welfare reforms are now accessible and scalable for modern husbandry and slaughter practices. Despite these significant improvements, contemporary practices, exemplified by factory farms, run afoul of the basic principles of good ILS rules. Nonetheless, certification programs, new regulations, and changes in attitudes pave the way for the approximation of a vision of humane sustainability.

Utopian visions involve radical change from contemporary visions and motivate greater long-term progress by "shifting the goal posts of common morality," enabling large-scale shifts in people's perceptions of animals and their moral significance. Utopian visions facilitate the gradual internalization of new rules that better reflect the verdicts of critical thinking. They "give concrete expression to the ideals of utilitarianism" and can make "achievable what once were utopian proposals" (Varner 2012: 256). Utopian visions are not implementable in the near term (because the principle of utility has a built-in "conservative bent or inertia"), but future generations will update the ILS rules of their ancestors, and by approximating utopian visions "in their personal moralities," individuals "can help move the goal posts of common morality and make their implementation in future generations more readily justifiable from a Harean perspective" (ibid. 260).

According to Varner, we should disallow the killing of animals, not even in utopian visions, at least now. Animal activists, he writes, "live a lifestyle that it would be unethical to force upon society at large at present, but which might become feasible as social norms in future generations… . In the long run … [they] can influence entrenched cultural assumptions and, over time, reduce the transition costs of reform" (ibid. 270). For now, "a reasonable case can be made for eating small quantities of meat from extensively reared ruminants (e.g., cattle) as a form of demi-vegetarianism"

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

(ibid. 283). Varner offers an illustration with Ernest Callenbach's (1996) "Buffalo Commons" proposal:[6]

> that dwindling populations in rural areas of the Great Plains be supported by a combination of electric wind farms and bison ranching on a truly grand scale … [N]ative prairie grasses are well adapted to the natural grazing habits of bison … This causes bison on "suitably large unfenced ranges" to naturally achieve the effects of Allen Savory's highly management-intensive (and fence-intensive) rotational grazing system for cattle (Varner 2012: 274; references omitted)

Bison being merely sentient, Varner finds the proposal "highly appealing aesthetically" and suggests that "bison rearing and slaughter could be made tremendously humane, given their natural behavior and resilient health, and modern carcass-processing technologies" (i.e., *in-situ* selective slaughter, with minimal pain, stress or herd disruption) (ibid.). Even though bison meat is still is a "niche market," the proposal is utopian in its "grand scope … a vision of humane sustainability that can be seen as inspiring to contemporary consumers and producers. Just describing it goes some small way toward moving the goal posts of our common morality" (ibid. 275). In fact, Wild Buffalo Ranchers in South Dakota "claim to do just this."

Varner's framework remains succinct, so let me make some extra assumptions. The difference between contemporary and utopian visions is *a matter of time and perspective*. The benefits of utopian visions are to be assessed in the long run, in part on the basis of their aspirational effects. Further, visions focus not just on individual behavior but on the *factors causing or facilitating behavior*. They do so through ILS rules and the social, political, and institutional infrastructures that support them. Finally, the process is *mutually reinforcing*: new attitudes and behavior support higher-level changes, higher-level changes shape attitudes and behavior, and behavior signals and reinforces attitudes. On these assumptions, a vision aimed at producing good ILS rules should include a timeline and a specification of its targets (legislation, public policy, moral education, consumer base, social media, etc.). I will argue that Varner's utopian visions may inadvertently foster the type of harmful attitudes that they are designed to suppress. In a nutshell: the permissible use of sentient creatures for food sets the bar too low for respect (i.e., promoting animals' interests for their own sake). As I argue now, Varner's visions miss important features of their target audiences.

The argument has two steps. First, utopian visions require a certain connection to contemporary visions in order to shift the goal posts. I briefly lay out the desiderata of an effective utopia – which provides motivation for greater change rather than entrenching the status quo. Second, because the promotion of visions has opportunity costs, a case for specific visions requires comparisons to alternatives, and I suggest that Varner's visions do not come out ahead on this front.

[6]Varner considers a third utopian proposal, besides demi-vegetarianism and the Buffalo Commons. Cellular agriculture and genetic engineering could be used to dispense with animals and/or eliminate their capacity for suffering (Varner 2012. 275–79). Yet, besides the obstacle of changing consumer preferences, Varner falls short of fully endorsing cellular agriculture since, by reducing the number of animals, it might decrease total happiness.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

4.2. Shifting the goal posts

The function of utopian visions of humane sustainability is to *inspire* contemporary consumers and producers. For various reasons, we can assume that truly humane agriculture would not work well at scale (by its own ethical standards) (Schlottmann and Sebo 2019: 108–10). The question is what trickle-down effects utopian visions can have on contemporary visions. The key, Varner and I agree, is encouraging a dramatic reduction in consumption of meat and animal products. Can humane agriculture deliver on its promise?

Let's suppose that visions like the Buffalo Commons would bring about significant benefits – by restoring grasslands, increasing biodiversity, enriching soil quality, and enhancing carbon sequestration; by giving pleasant lives to animals that could be harvested humanely with mobile processing plants and efficient slaughter techniques; and by enabling sustainable agricultural practices that benefit farmers and local communities. Using the same land for other purposes, whether it be untouched grasslands, less sustainable extensive ranching, or growing crops, would not produce more marginal utility. Let's concede too that, if they were enforced by archangels, such visions would maximize total happiness. Even so, the comparative benefits of different rules depend on the consequences of internalizing and enforcing them. In other words, Varner must show that associating humane sustainability with the continuation of using animals as sources of food and ecological services at the intuitive level will most likely shift goal posts in the right direction. For utopian visions to be effective, and worth striving for, they must be such that, if enough people were to internalize them, they would actually improve animals' lives. But, as Varner acknowledges throughout, these are complex empirical issues.

Something the Harean wants to avoid is unintended spillover: for instance, that embedding replaceability in utopian visions could have a widespread negative effect on contemporary visions by entrenching the status quo, where eating animals is seen as morally acceptable. And because contemporary visions involve *billions* of animals, the stakes are high. Critical thinking, ruled by utility, can act as a safeguard, but it cannot do the work alone. We must ensure that ILS rules will reliably implement the correct verdicts of critical thinking. Under these conditions, why not think that animal liberation, or any relevant approximation, could have the desirable effect on common morality by challenging the status quo?

4.3. Opportunity costs

Humane agriculture, by promoting niche markets (e.g. grass-fed beef, pasture-raised pork, and cage-free eggs), does not involve very many animals. So whatever harms it still causes to animals should not be cause for concern relative to their overall positive effects. And, from replaceability, we can infer that we are doing more good than harm by giving these animals happy lives. Plus, more people become sensitive to animal welfare, and laws and common morality will shift accordingly. Correct?

Not so fast. Promoting utopian visions has opportunity costs. Time spent promoting humane agriculture is not spent promoting plant-based alternatives. So even if humane utopias did not have a negative effect on animals over time, alternatives might still have better consequences. And contrary to what Hare (1999: 241) thought, demi-vegetarians may not have a "very powerful advantage over full vegetarians when it comes to influencing the market." Farmers are rewarded by converting primary resources and labor

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

into animal protein efficiently. Unless it is economically beneficial for them to do so, they have no incentives to take animals' interest directly into account. Welfare standards are primarily a response to demand from well-off consumers and constrained by economic interests more than moral considerations. A Harean might reply that humane sustainability would put constraints on economic trade-offs since animal welfare is a fundamental desideratum of humane sustainability. But the initial question recurs: can demi-vegetarianism and humane agriculture, as utopian visions, deliver on their promise of shifting the goal posts more effectively than alternatives?

First consider what it would take for alternatives to gain traction. Meat is so inexpensive because the production and distribution costs are kept low by economies of scale, efficiencies, government support (subsidies, checkoffs, and buybacks), and unregulated cost externalization (Bollard 2019). One way to make alternatives competitive is to increase their consumer base and for new interest groups to organize. Producing scalable, sustainable plant-based alternatives could re-anchor groups to new social norms much more effectively than humane agriculture. Goal post-shifting may involve disrupting the global food system rather than promoting small-scale animal husbandry.[7]

Still, we can treat what visions would be most effective at inspiring widespread change as an open question, especially if we remember that utopian visions involve a timeline and moving targets. We should distinguish between *advocating* and *adopting* visions. Generally, campaign A might cause more people to decrease their consumption of animal products, and by greater margins, than Campaign B, independently of whether we should actually adopt the diet recommended by A or B. Thus, Varner cites "certain statistics" suggesting "that Hare may be right," that people are not ready to be converted to full vegetarianism (2012: 272). Yet Varner acknowledges that "which has greater 'propaganda value' – demi-vegetarianism or veganism – is an empirical [question]."

Indeed, but there are some empirical reasons to think Hare may not have been right. Recall that ILS rules must be relatively inflexible to be efficient. So, when we make ethical consumption choices, we must also think about how they will affect our habits, which will in turn affect our future actions, and so on. Schlottmann and Sebo (2019: 189) write: "one reason … to take a relatively strict approach to ethical consumption … is that anything less than that makes it difficult to develop the relevant habits, and so to live up to the relevant ideals in the hustle and bustle of everyday life." To conscientious omnivores and demi-vegetarians, we might say: "if we allow ourselves to eat meat at all, then we will likely end up eating meat more often than we should" (ibid.). This slippery slope argument is behavioral and might turn out to be incorrect, but if we can already personally suspect that we would act rightly less often if we adopted a flexible diet, then we know we should internalize stricter habits. After all, the very reason we need intuitive rules is that the more complexity we introduce in our decision procedures, the less likely we are to deploy those rules reliably. The finer-grained the rule, the more room for error, complacency or rationalization (also see John and Sebo, forthcoming). Indeed, Rothgerber (2015) found evidence that conscientious omnivores are less likely than vegetarians to believe they should strictly adhere to their diet, and partly as a result are more likely to violate their diet and feel less guilt when doing so. One explanation seems to be "difficulty following one's diet." This might be surprising since conscientious omnivores "have more options than meat

[7]See e.g. the Good Food Institute at http://www.gfi.org (accessed on September 5, 2019).

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

abstainers" (ibid. 202), but it makes sense if finer-grained (i.e., more complex rules) are harder to abide by.

Moreover, flexible ethical consumption lends itself to *moral licensing*, "the tendency to see what one is currently doing as all one needs to be doing in order to be a good person" (Schlottmann and Sebo 2019: 190). I suggested that demi-vegetarianism may further entrench the status quo of perceiving farm animals as edible. Moral licensing likely contributes to that process too. Demi-vegetarianism looks less demanding than a full commitment to conscientious omnivory, but by the same token it lends itself to complacency, ultimately acting as a drag on moral change.

We must also consider the *social context* of our choices – our friends, family, colleagues, and the businesses we patronize. We influence them and they influence us. In social contexts, compliance signals approval of a desirable social norm (Bicchieri 2017). Legality and market availability also signal social acceptability.[8] So, if most people eat meat, even in small quantities, most people will uptake the norm that eating animals is acceptable (Delon 2018; Schlottmann and Sebo 2019). The smaller the gap between utopian visions and what is relatively easy under contemporary visions but still counts as progress, the smaller the effect of utopian visions. Of course, too *great* a gap could be paralyzing too, making meaningful progress seem unattainable. Still, this leaves plenty of leeway. Since our behavior is largely a function of social norms, behavioral change requires the norm to shift, which, in turn, requires that enough people be ready to change their behavior for a new norm to emerge (Bicchieri 2017). So, we must choose policies that we are disposed to follow *and* that make us less likely to cause harm. The above empirical speculations suggest that alternative visions like animal liberation may hold more promise for inspiring change.[9]

This doesn't mean we cannot meet in the middle. By focusing exclusively on short-term gains, incremental welfare campaigns (e.g., cage-free eggs and Meatless Monday) yield immediate benefits, although they may fail to sustain profound, durable shifts in attitudes. Off-putting, confrontational tactics, direct action, and abolitionist messaging have immediate costs, although the long-term benefits of promoting a society built upon compassion rather than exploitation and killing may outweigh such costs. We would do best to see these different tactics as mutually reinforcing rather than exclusive (Sebo and Singer 2018), as Varner eventually does. He considers Singer's remark that, even if raising animals for food may be justifiable critically, "the level of practical moral principles" (i.e., ILS rules) is different:

> Killing animals for food makes us think of them as objects that we can use as we please … How can we encourage people to respect animals, and have equal concern for their interests, if they continue to eat them for their mere enjoyment? To foster the right attitudes of consideration for animals, including non-self-conscious ones, it may be best to make it a simple principle to avoid killing them for food. (Singer 2011: 134 [1993 edn])

[8]Legal scholars have argued that the law generates compliance through its expressive function, by signaling the acceptability of the very norms it seeks to enforce, rather than merely (or primarily) through coercion (Bilz and Nadler 2014; McAdams 2015).

[9]For instance, Gruen and Jones (2015) argue for "veganism as an aspiration," expressing both the inevitability of causing harm in our daily lives and the ethical commitment to regard sentient animals as non-edible. It was brought to my attention while revising this article that John and Sebo (forthcoming) argue from an indirect consequentialism perspective against the Logic of the Larder.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

Singer's position is actually more nuanced, acknowledging that different audiences will be more receptive to different diets and lifestyles and that "[i]t is probably a good thing if different people and organizations promote each of these goals" (Singer 1999: 325; Varner 2012: 273). Again, Varner treats this point as an open question. He seems confident that we can draw fine-grained distinctions among animals that are good to eat and those that are not, as he is that we can compartmentalize the intuitive and critical levels (more on this shortly). I'll argue it's not quite clear.

Still, one could contemplate different visions concurrently across different audiences, or across time, moderate visions making people more favorably disposed toward radical ones over time. If these visions are mutually reinforcing rather than exclusive (one utopia being nearer on the timeline than the other), the transition would nonetheless require that we already start promoting the more radical vision, at least to some audiences. I suspect Varner might agree. But this doesn't settle the case yet.

Let's recap first. What utopian visions we should promote (i.e., advocate rather than adopt) depends on a range of empirical considerations, including a timeline and a specification of the target audiences of particular sets of ILS rules.[10] But if the above speculations are correct, we should design ILS rules guided by sufficiently radical utopian visions, lest we fail to deliver on the promise of humane sustainability. How exactly to articulate such utopian visions so as not to hinder change in the here and now is an open question, but, at any rate, our end-point should not be humane agriculture.

In the next section, I argue that TLU falls prey to compartmentalization problems.

5. Problems of compartmentalization

By way of preface, let me unpack three apparent assumptions made by the claim that humane slaughter-based agriculture would be optimific. They operate at different levels:

1. Killing sentient animals does not cause them morally significant all-considered harm. [Critical and intuitive levels][11]
2. Killing sentient animals is not perceived as harmful to them by ordinary human beings. [Intuitive level]
3. How we should treat animals ultimately depends on circumstances, including consumer preferences, because it depends on the total value of outcomes rather than the value of individuals. [Critical level]

(1) is debated even among utilitarians (Višak and Garner 2016). (2) is an empirical assumption. Psychology suggests that people do not consider killing food animals to be morally *wrong*, yet many people disapprove of killings they perceive as unnecessary, presumably out of concern for the animals themselves and because they see killing as nonetheless *harmful* (Bastian and Loughnan 2017). (3) is characteristically utilitarian. At the intuitive level, (3) seems at odds with promoting *respect* for all sentient animals. So it must remain sealed at the critical level. But it's only with (3) that (1) and (2) can do the work for ILS rules. On the other hand, (3) would be innocuous if we dropped (1)

[10]Research is conducted by Animal Charity Evaluators, Faunalytics, and The Humane League Labs, among others. My claim is simply that there remains much empirical uncertainty.

[11]If the merely sentient lack the capacities to ground an interest in life, they are *not* denied due consideration when they are killed. However, TLU fosters a conception of the merely sentient as having lesser moral status than they have on Varner's own account.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

or (2). If killing sentient animals were wrong, or perceived to be so as a result of internalizing ILS rules, the effects of (3) on animals would likely not be too detrimental (assuming standards of welfare are maintained). But if we hold all three assumptions, TLU is asking consumers to regard sentient creatures as worthy of respect while justifying their use on the basis of its impersonal benefits. Varner thus builds not just a specific view of the harm of death but a consequentialist axiology into ILS rules. Most people do not intuitively conceive of the value of lives in such terms, which leads to two problems of compartmentalization.

5.1. Williams' objection to Hare

I argue that Varner's conception of the value of lives conflicts with a standard understanding of moral status and the cognitive templates that constrain good ILS rules. This is not to say that no agent could possibly be equipped with the relevant psychology, but the point of designing descriptively adequate ILS rules is to accommodate us more or less as we are.

Let's begin by reviewing Bernard Williams' objection to Hare and Varner's response to the problem – or problems – of compartmentalization. According to Williams, Hare's TLU involves an incoherent moral psychology. It's important for TLU to withstand the objection for the sake of descriptive adequacy. Hare anticipated the problem:

> To say that it is impossible to keep critical and intuitive thinking going in the same thought-process is like saying that in a battle a commander cannot at the same time be thinking of the details of tactics, the overall aim of victory, and the principles … which he has learnt when learning his trade. (1981: 52)

This is what Williams suggests is impossible. Williams criticized Hare in two main instances (cf. Miller 2014).

> The [indirect utilitarian] dispositions are seen as devices for generating certain actions, and those actions are the means by which certain states of affairs, yielding the most welfare, come about. This is what these dispositions look like when seen from outside, from the point of view of the utilitarian consciousness. But it is not what they seem like from the inside. Indeed, the utilitarian argument implies that they should *not* seem like that from the inside. The dispositions help to form the character of an agent who has them, and they will do the job the theory has given them only if the agent does not see his character purely instrumentally … (Williams 1985: 107–08)

Williams is making two related points. First, dispositions (i.e., prima facie principles, ILS rules) are seen from a certain perspective, which affects their character. And since these dispositions are designed to be effective, they cannot do their job if they are to be seen or experienced from two incompatible perspectives. Second, their effectiveness turns on our intuitive perception of certain things as intrinsically valuable, which, again, the act-utilitarian standpoint undermines. This point, I think, straightforwardly applies to our perception of the value of animal lives.

In the second instance, Williams wrote that

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

> [Hare's theory] represents the intuitive responses as deeply entrenched, surrounded by strong moral emotions, sufficiently robust to see the agent through situations in which sophisticated reflection might lead him astray, and so on; and yet at the same time explains those responses as a device to secure utilitarian outcomes. The theory ignores the fact that the responses are not merely a black-box mechanism to generate what is probably the best outcome under confusing conditions. Rather, they constitute a way of seeing the situation, and you cannot combine seeing the situation in that way, *from* the point of view of those dispositions, with seeing it in the archangel's way … (Williams 1988: 189–90)

Williams' critique is internal: ILS rules cannot do their job because they're chosen using critical thinking (from the lens of maximizing aggregate utility). The motivational force of rules depends on their shaping the way we see the practical world. The key step in Williams' argument is that we cannot see something both as intrinsically valuable and as instrumentally necessary (aggregation-wise). When we see rules as mere instruments, the objection goes, we are not using them in the way they were designed to be used – we can *no longer* do so – yet this is precisely what TLU implies for critical thinking to play its role. By treating animals as edible commodities, we threaten the resilience of intuitive dispositions by irreversibly introducing market norms in the moral domain (see e.g. Ariely 2008: 76). Williams did not think that it was "impossibly difficult" to recover from the utilitarian perspective, but "that the thoughts are not stable under reflection" (1988: 190).

Both Hare and Varner feel capable of shifting between the two perspectives, and keeping them distinct as required (Hare 1981: 52; Hare 1988: 289–90; Varner 2012: 84). Varner adduces some empirical evidence that many people are capable of what Williams thought a true two-level utilitarian would have to possess: not just gut reactions evolved by our ancestors, but dispositions that would have to be both motivationally significant and open to critical inquiry.

The main piece of evidence comes from dual-process theories, but it's unclear whether it supports TLU. The literature does show that people rely on two distinct types of thinking, but our intuitions are not those that would have been designed by critical thinking. If anything, dual-process theorists emphasize the unreliability of "automatic settings," or system 1. Our intuitions are likely to be unreliable because they are rather inflexible (Greene 2013: 217). For instance, they are insensitive to scope and susceptible to "modular myopia": "our brains have a cognitive subsystem … that monitors our behavioral plans and sounds an emotional alarm bell when we contemplate harming other people… . [T]his alarm system is '*myopic*,' because it is *blind to harmful side effects*" (ibid. 224). Our intuitions can be unreliable, even if they work well in many contexts, because "[a]ll automatic settings rely on specific cues that are only imperfectly related to the things they're designed to detect" (ibid. 227).

Here's the worry: TLU seeks to design efficient ILS rules that will, by their very nature, be "rather inflexible." Given the goal of fostering moral consideration for sentient animals, they will likely have to rely on harm-based cues. But precisely insofar as they are inflexible, there are going to be contexts where our harm-detecting module may clash with the permissibility of eating animals. This is, presumably, where Greene's *manual* mode would kick in, to override the automatic settings and lay out the utilitarian rationale for replaceability. We are, after all, "emotionally blind" but not "cognitively blind." So far, so good. But if the literature on cognitive dissonance is any indication, as

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

I discuss shortly, we should expect this process to be quite uncomfortable, which Greene would be the first to concede. The challenge is for that discomfort not to backfire and undermine the rather inflexible moral consideration owed to sentient animals according to ILS rules.

Dale Miller argues that our reactive attitudes are resilient enough to survive access from the objective standpoint. We can't "avoid experiencing [them] in the ordinary way the vast majority of the time. We can displace these attitudes with an objective attitude by choice, as in Hare's cool hour, but not for long" (2014: 54). Likewise, Miller suggests, we can alternate between seemingly incompatible types of moral thinking over time. Williams would have to show that a TLU agent "would *always* think critically and intuitively simultaneously" (ibid. 55). But, Miller and Varner claim, we can compartmentalize our perspectives to preserve the viability of each system.

I agree that reactive attitudes are resilient. Perhaps Williams was wrong that our intuitive dispositions would not survive utilitarian thinking. Still, the following must be true to dissolve the problem of compartmentalization. First, dual-process theories must be correct – in itself a strong assumption[12] – so that system 2 can override system 1. Second, intuitive thinking must be resilient enough that system 2 does not prevent us from resuming thinking intuitively when we need to. Of course, we should not trust system 1 to tell us when we should do so, and so we're left to wonder how system 2 can lead us to resume thinking intuitively. Yes, critical thinking designs ILS rules, but can it tell us *how* to shift gears? There must be a sweet spot of rules that are resilient enough to fulfill their role yet flexible enough that critical thinking can control them. This is what a descriptively adequate TLU seems committed to. Respect-cum-replaceability must be among such rules. Section 5.2 will cast doubt on this possibility.

5.2. Compartmentalizing animals

On one influential model, moral cognition consists in the perception of two minds. Roughly, people attribute moral properties on the basis of *mind perception*. Attributions of moral standing track perceptions of "mindedness," typically along two dimensions: Experience (e.g. suffering, feelings, emotions) and/or Agency (e.g. thinking, memory, planning, self-control) (Gray et al. 2007). Morally significant events involve a "moral dyad," a pairing of a (blameworthy) agent and a (vulnerable) patient: say, a villain torturing a kitten; a hero rescuing an old lady. The moral dyad, with harm at its core, constitutes the fundamental template for perceiving moral violations (Gray et al. 2012; Gray et al. 2014).

A growing psychological literature has established the role of mind perception – and factors and biases that influence it – in our attitudes and behavior toward animals (Goodwin 2015; Jack and Robbins 2012; Kasperbauer 2017; Sytsma and Machery 2012). The "meat paradox" is the well-documented fact that most people (claim to) believe it is morally wrong to harm sentient creatures yet (believe that it is permissible to) eat animals. As a result, meat-eaters tend to experience cognitive dissonance when holding both attitudes. In order to reduce the discomfort, they have two options: "one can reject meat consumption, bringing one's behaviors into alignment with one's moral ideals, or one can bring one's beliefs and attitudes in line with one's behavior through various psychological maneuvers" (Piazza et al. 2015: 114). *Motivated cognition* explains

[12]Gray et al. (2014), Railton (2014) and Woodward (2016) cast doubt on the affective/rational divide.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

much of it (Piazza and Loughnan 2016). People can deny that animals have the psychological features that ground their moral status, typically the capacity to suffer: Bastian et al. (2012) indicates a negative relation between the perceived edibility of an animal and how "minded" it is perceived to be. People can also deny that animals deserve moral concern to begin with: Loughnan et al. (2010) primed subjects with either dried beef or dried nuts and found that the former leads to decreased attribution of a mind and moral status to cows. Bratanova et al. (2011) found that when some animals are categorized as food (e.g., Papua New Guinea's tree kangaroos), participants attribute to them less sentience and lesser moral status. The fact that people engage in motivated cognition in order to reduce dissonance confirms the tendency to attribute a mind, and consequently moral standing, to other animals. A series of studies also found that omnivores are motivated to adopt a range of rationalizations for eating meat, including the so-called "4Ns": that it is natural, normal, necessary, and nice (i.e., pleasant) (Piazza et al. 2015).

These facts – the centrality of harm-based concern, the tendency to motivated reasoning – are central to the way we relate to other animals and might undermine the prospects of respect-cum-replaceability, because replaceability leads us to regard animals as deserving less consideration than they do, whereas respect leads us to disapprove of harming animals. And our acceptance of replaceability is all too likely to trigger complacent motivated reasoning. If these facts constrain ordinary psychology, ILS rules should be mindful of them.

Varner is aware of the pitfalls of building too much utilitarian thinking into ILS rules. Recently, he has argued that ILS rules "the internalization of which disposes one *not* to think like a utilitarian during the moment-to-moment management of many interpersonal relationships" should apply to our relationships with our pets (Varner 2017: 77). Because of the overall benefits of positive relationships with pets, he argues, pet-keepers should not see them as replaceable. For "loving another sentient being requires you not to think of it as 'replaceable' the way everyone and everything is when you are doing explicitly utilitarian thinking" (ibid.). The ILS rules governing our relationships with farm animals are notably different. This aligns with our ability to compartmentalize animals between different categories (Amiot et al. 2019). But, if Varner is correct about pets, the replaceability of farm animals precludes a range of attitudes that foster mutually beneficial relationships with them (of the sort that conscientious omnivores commonly claim to promote). In fact, why not imagine that Varner's (2002) conception of "domesticated partnerships" could provide utopian visions for *all* domesticated animals? Like companion animals, farm animals could generate revenue (through tourism, therapy, veterinarian care, or labor) in non-exploitative relationships that would exercise their faculties. Since Varner assumes that we should not regard our pets as replaceable, let alone eat them, it's unclear why we really need different ILS rules for different categories of animals, especially if the benefits of small-scale humane slaughter-based agriculture are bound to be so restricted by scale.[13]

[13]Cats, unlike dogs, are considered "obligate carnivores," and so generally cannot be fed a plant-based diet. Their lives thus depend on the by-products of factory farming (Milburn 2017). I tend to think of this as a later generation, if important, problem, requiring adequate utopian visions and technology. I suspect that "veganism as an aspiration" (Gruen and Jones 2015) might provide incentives among consumers and producers to gradually shift norms until humane sustainable vegetarian diets are available for pets. Thanks to Dale Miller for raising the issue.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

Admittedly, in *some* sense replaceability is consistent with respect. As many cat and dog lovers will attest to, we regard our companions as worthy of respect, we provide them with individualized care, yet when they die, it's not infrequent for us to "replace" them in some sense. How does this differ from the sense of replaceability implied in the replaceability argument? For one thing, we mourn our deceased companions, remember them qua individuals, and don't so much try to replace them as we seek other companions to fill their role. They are never truly replaced, which is why we grieve. On the other hand, animal farming treats animals as strictly replaceable: the primary reason for bringing new beings into existence is that they, in turn, will be replaced by other animals in the context of market exchanges. Such interactions are closed once the animal dies, and parties are accordingly substitutable. This is what market norms allow for, through impersonal relationships (Maguire and Brown 2019). Markets would cease to be as efficient if farmers developed the sort of non-substitutable relationships that genuine respect requires. Market efficiency is a major benefit of impersonal transactions (ibid. 10–13). In *this* sense, replaceability is hardly compatible with respect, if respect implies a complex of personal, open-ended attitudes and dispositions.

Couldn't ranchers internalize certain ILS rules such that they would want to care for bison for their own sake, not just for profit? They would have them humanely slaughtered while sincerely believing that death does not violate any morally significant interests, since the merely sentient have no interest in continued existence. They would be respecting bison while treating them as replaceable. This is plausible, and remember that ILS rules can be tailored to variations in "personal morality." Yet I doubt that this can be generalized at the intuitive level – especially among consumers, who do not stand in any meaningful relationships with farm animals and treat them, as a matter of fact, as commodities. If it turns out that more people than I suspect can in fact respect animals while regarding them as replaceable, then this particular compartmentalization problem dissipates. I would welcome a two-level utilitarian argument to this effect. But Varner himself doubts this is feasible with pets; it remains unclear why this would be much easier with farm animals. We haven't fully solved, I conclude, the problems of compartmentalization.

6. Conclusion

Varner provides a forceful defense of humane agriculture and demi-vegetarianism, but I argued that TLU can explain why it fails to value humane lives appropriately as well as how it can be fixed. I began by unpacking key assumptions of Varner's view, including the commitment to replaceability. I then argued that his visions framework can generate conclusions that support alternative utopian visions rather than humane slaughter-based agriculture. Finally, I argued that TLU is susceptible to problems of compartmentalization that undermine the prospects of accommodating our duty to respect animals while regarding them as edible. ILS rules promote attitudes that are likely to lead people to violate the very standards of welfare that humane sustainability upholds are not optimific by TLU's own lights. Fortunately, TLU also illuminates alternatives that do not fall prey to similar problems.[14]

[14]Thanks for valuable feedback to audiences at the 2016 meeting of the International Society for Environmental Ethics and the 2016 Rocky Mountain Ethics Congress, Duncan Purves, Jeff Sebo, Dale Miller, and two anonymous referees for *Utilitas*.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

References

Amiot C. E., O. Lépine, and J. Zaky. 2019. Compartmentalization of animals: toward an understanding of how we create cognitive distinctions between animals and their implications, *Analyses of Social Issues and Public Policy*, Early View: 1–27. https://doi.org/10.1111/asap.12177.

Ariely D. 2008. *Predictably Irrational* (New York: HarperCollins)

Bastian B., S. Loughnan, N. Haslam, and H. R. M. Radke. 2012. Don't mind meat? The denial of mind to animals used for human consumption, *Personality and Social Psychology Bulletin*, **38**(2): 247–56.

Bastian B., and S. Loughnan. 2017. Resolving the meat-paradox: a motivational account of morally troublesome behavior and its maintenance, *Personality and Social Psychology Review*, **21**(3): 278–99.

Bentham J. 1907 [1789]. *Introduction to the Principles of Morals and Legislation* (Oxford: Clarendon Press).

Bicchieri C. 2017. *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms* (Oxford: Oxford University Press).

Bilz K., and J. Nadler. 2014. Law, moral attitudes, and behavioral change, in *The Oxford Handbook of Behavioral Economics and the Law*, ed. by E. Zamir and D. Teichman (Oxford: Oxford University Press), pp. 241–67.

Bollard L. 2019. Why is meat so cheap? *The Open Philanthropy Project Farm Animal Welfare Newsletter*. October 4, 2019. https://mailchi.mp/1a041ae10c0d/why-is-meat-so-cheap?e=af5349800b. Accessed November 7, 2019.

Bratanova B., S. Loughnan, and B. Bastian. 2011. The effect of categorization as food on the perceived moral standing of animals, *Appetite*, **57**: 193–96.

Callenbach E. 1996. *Bring Back the Buffalo! A Sustainable Future for America's Great Plains* (Berkeley: University of California Press).

Delon N. 2016. The replaceability argument in the ethics of animal husbandry, in *Encyclopedia of Food and Agricultural Ethics* (2nd edn), ed. by P. B. Thompson and D. M. Kaplan (Dordrecht: Springer). https://doi.org/10.1007/978-94-007-6167-4.

Delon N. 2018. Social norms and farm animal protection. *Palgrave Communications*, **4**(139). https://doi.org/10.1057/s41599-018-0194-5.

Goodwin G. 2015. Experimental approaches to moral standing, *Philosophy Compass*, **10/11**: 914–26.

Gray H., K. Gray, and D. M. Wegner. 2007. Dimensions of mind perception, *Science*, **315**: 619.

Gray K., C. Schein, and A. F. Ward. 2014. The myth of harmless wrongs in moral cognition, *The Journal of Experimental Psychology: General*, **143**: 1600–15.

Gray K., L. Young, and A. Waytz. 2012. Mind perception is the essence of morality, *Psychological Inquiry*, **23**(2): 101–24.

Greene J. 2013. *Moral Tribes: Emotion, Reason, and the Gap between Us and Them* (New York: Penguin).

Gruen L., and R. C. Jones. 2015. Veganism as an aspiration, in *The Moral Complexities of Eating Meat*, ed. by B. Bramble and B. Fischer (Oxford: Oxford University Press), pp. 153–71.

Hare R. M. 1981. *Moral Thinking: Its Levels, Method and Point* (Oxford: Oxford University Press).

Hare R. M. 1988. Comments on Williams, in *Hare and Critics: Essays on Moral Thinking*, ed. by D. Seanor and N. Fotion (Oxford: Clarendon Press).

Hare R. M. 1993. Possible People, in *Essays on Bioethics* (Oxford: Clarendon Press), pp. 67–83.

Hare R. M. 1999. Why I am only a demi-vegetarian, in *Singer and his Critics*, ed. by D. Jamieson (Oxford: Blackwell), pp. 233–46.

Jack A. I., and P. Robbins. 2012. The phenomenal stance revisited, *Review of Philosophy & Psychology*, **3** (3): 383–403.

John T. M., and J. Sebo. Forthcoming. Consequentialism and nonhuman animals, in *The Oxford Handbook of Consequentialism*, ed. by D. Portmore (Oxford: Oxford University Press).

Kahneman D. 2011. *Thinking, Fast and Slow* (New York: Farrar, Strauss and Giroux).

Kasperbauer T. J. 2017. Mentalizing animals: implications for moral psychology and animal ethics, *Philosophical Studies*, **174**(2): 465–84.

de Lazari-Radek K., and P. Singer. 2014. *The Point of View of the Universe* (Oxford: Oxford University Press).

Levy S. S. 1994. The coherence of two-level utilitarianism: Hare vs. Williams, *Utilitas*, **6**(2): 301–09.

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press

Loughnan S., N. Haslam, and B. Bastian. 2010. The role of meat consumption in the denial of moral status and mind to meat animals, *Appetite*, **55**: 156–59.

Maguire B., and B. Brown. 2019. Markets, interpersonal practices, and signal distortion, *Philosophers' Imprint*, **19**(14): 1–16.

McAdams R. H. 2015. *The Expressive Powers of Law: Theories and Limits* (Cambridge, MA: Harvard University Press).

Milburn J. 2017. The animal lovers' paradox? On the ethics of "pet food", in *Pets and People: The Ethics of our Relationships with Companion Animals*, ed. by C. Overall (Oxford: Oxford University Press), pp. 187–202.

Miller D. E. 2014. Reactive attitudes and the Hare–Williams debate: towards a new consequentialist moral psychology, *The Philosophical Quarterly*, **64**(254): 39–59.

Piazza J., M. B. Ruby, S. Loughnan, M. Luong, J. Kulik, H. M. Watkins, and M. Seigerman. 2015. Rationalizing meat consumption. The 4Ns, *Appetite*, **91**: 114–28.

Piazza J., and S. Loughnan. 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): 867–74.

Railton P. 2014. The affective dog and its rational tale: intuition and attunement, *Ethics*, **124**(4): 813–59.

Rothgerber H. 2015. Can you have your meat and eat it too? Conscientious omnivores, vegetarians, and adherence to diet, *Appetite*, **84**: 196–203.

Salt H. S. 1914. The Logic of the Larder, in *The Humanities of Diet* (Manchester: The Vegetarian Society).

Sebo J., and P. Singer. 2018. Activism, in *Critical Terms for Animal Studies*, ed. by L. Gruen (The University of Chicago Press), pp. 33–46.

Schlottmann C., and J. Sebo. 2019. *Food, Animals, and the Environment: An Ethical Approach* (New York: Routledge).

Singer P. 1999. A response, in *Singer and his Critics*, ed. by D. Jamieson (Malden: Blackwell), pp. 269–335.

Singer P. 2011. *Practical Ethics*, 3rd edn (Cambridge: Cambridge University Press) [First published in 1993.]

Singer P. 2016. Afterword, in *The Ethics of Killing Animals*, ed. by T. Višak and R. Garner (Oxford: Oxford University Press), pp. 229–35.

Stephen L. 1896. *Social Rights and Duties*, vol. 1 (London: Sonnenschein).

Sytsma J., and E. Machery. 2012. The two sources of moral standing, *Review of Philosophy & Psychology*, **3**(3): 303–24.

Tulving E. 1985. Memory and consciousness, *Canadian Psychology*, **26**: 1–12.

Varner G. 2002. Pets, companion animals, and domesticated partners, in *Ethics for Everyday*, ed. by D. Benatar (New York: McGraw-Hill), pp. 450–75.

Varner G. 2012. *Personhood, Ethics, and Animal Cognition* (Oxford: Oxford University Press).

Varner G. 2017. A two-level utilitarian analysis of relationships with pets, in *Pets and People: The Ethics of our Relationships with Companion Animals*, ed. by C. Overall (Oxford: Oxford University Press), pp. 64–79.

Višak T. 2016. Do utilitarians need to accept the replaceability argument?, in *The Ethics of Killing Animals*, ed. by T. Višak and R. Garner (Oxford: Oxford University Press), pp. 117–35.

Višak T., and R. Garner (eds). 2016. *The Ethics of Killing Animals* (Oxford: Oxford University Press).

Williams B. 1985. *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press).

Williams B. 1988. The structure of Hare's theory, in *Hare and Critics: Essays on Moral Thinking*, ed. by D. Seanor and N. Fotion (Oxford: Clarendon Press), pp. 185–96.

Wolfson D. J., and M. Sullivan. 2004. Foxes in the hen house: animals, agribusiness, and the law: a modern American fable, in *Animal Rights: Current Debates and New Directions*, ed. by C. R. Sunstein and M. C. Nussbaum (Oxford: Oxford University Press), pp. 205–33.

Woodward J. 2016. Emotion versus cognition in moral decision-making: a dubious dichotomy, in *Moral Brains: The Neuroscience of Morality*, ed. by S. M. Liao (New York: Oxford University Press), pp. 87–116.

Cite this article: Delon N (2020). Valuing humane lives in two-level utilitarianism. *Utilitas* **32**, 276–293. https://doi.org/10.1017/S0953820819000542

https://doi.org/10.1017/S0953820819000542 Published online by Cambridge University Press