Shelly Kagan, *How to Count Animals More or Less.* Oxford: Oxford University Press, 2019. Pp. x, 309.

You've just been named to your university's animal research committee. Keen to prove to your colleagues from around the university that philosophy isn't just an idle wheel, you decide to brush up on the latest thinking on questions like: when, if ever, should we be willing to trade off animal welfare against expected future human welfare? Or perhaps you don't like thinking in terms of welfare. In that case: when, if ever, is it permissible for us to override whatever sorts of rights animals have in the interest of saving, or at least significantly improving, human lives? Either way, you decide to pick up Shelly Kagan's new book, nominally on just these topics.

I'm afraid that you are likely to be disappointed. The problem isn't one of false advertising exactly. Kagan is interested in just these questions. But he addresses them at such a high degree of abstraction that it's very difficult to see how the discussion is going to be of much help. To be clear, Kagan is admirably up-front about this. Indeed, he starts the book by telling us that "I suppose this does indeed count as a topic *within* practical ethics, but I fear that the discussion is about as abstract a treatment of the topic as one could offer" (ix). The last two sentences of the book strike a similar note: "for the most part all I have really done is to offer some rough, preliminary suggestions about what a suitable account might conceivably look like. The devil will be in the details, and the details are yet to come" (304). I take this to be an apt assessment of the state of things by the end of this work: some rough preliminary suggestions have been offered, but a great many important details are absent. I rather admire such a clear-headed assessment on the part of the author. Still, I found myself confused as to why neither the author nor the publisher viewed this as grounds for substantial revisions. This is a three-hundred page work, after all; one might have hoped for a few more of those details.

So what exactly does Kagan cover in the span of this work? Chapter 1 introduces some background notions like *moral status*, chapter 2 sketches Kagan's target, the 'unitarian' thesis that any creature worthy of moral consideration is worthy of equal moral consideration, and chapter 3 runs through various consequentialist distribution principles, like prioritarianism and sufficientarianism. Chapter 4 argues that any plausible view about welfare distribution will have to reject unitarianism and embrace what Kagan calls a 'hierarchical' view of moral status. Chapter 5 basically expands on Kagan (2016)'s suggestion that, whatever exactly grounds moral status, it involves a modal element. Chapter 6 considers some objections to the hierarchical view, and then chapters 7-11 basically run through the same dialectic again, this time assuming a deontological framework rather than a consequentialist one. I'll start with the consequentialist version of things.

Kagan's focus is on what he takes to be a quandary for consequentialists who care about animals. For simplicity's sake, let's call whatever good or combination of goods ought to be promoted 'welfare'. Now, let's further suppose that at least some animals are capable of having welfare, of having lives that can go better or worse in a morally-relevant sense. In our present world, there are myriad trade-offs between human welfare and the welfare of non-human animals. As good consequentialists, we should ask: when are these trade-offs justified? As a concrete case—though one that Kagan never actually considers—we might reflect on the ethical status of animal research. Even die-hard vegetarians may be inclined to think it morally permissible for a new vaccine to be tested

on laboratory mice before it is tested on human beings. If this inclination is to be justified, it looks like the vegetarian must be tacitly assuming that humans are worthy of *more* moral consideration than laboratory mice are. Alternatively, our vegetarian could revisit her initial assumption that lab mice are worthy of moral consideration at all. But Kagan is interested in how our moral theory should look on the assumption that (many) animals are worthy of at least some moral consideration.

In broad outline, Kagan's suggestion (developed in chapter 4) is that we should think of human beings as *counting for more* than e.g. lab mice. The difficulty comes in spelling out what that 'counting for more' amounts to. One option, which Kagan ultimately finds wanting, would be to claim that the average human life involves more welfare at any given moment than that of an average lab mouse. We have a capacity to feel more complex emotions than mice do, to form deeper loving bonds, to make complex plans both individually and together, etc. Human beings also tend to live much longer than mice. So whereas each mouse life we sacrifice involves *x* number of remaining good years, an average human life saved by the vaccine is likely to involve some multiple of this. As a rough-and-ready calculation of expected goodness, we can tally up average lab mouse welfare times average years sacrificed times the number of mousy lives required and weigh that against average human welfare times average years left to live times the number of humans who will benefit from the vaccine. Granted, this is only a rough-and-ready calculation; a great deal hinges on whether there are other, more promising vaccine-candidates out there, or whether climate change is likely to lead to mass human extinction in the near future.

Kagan takes this sort of reasoning to overlook another important way in which human beings 'count for more'. So far, we've basically treated human welfare and mouse welfare as equivalent, unit for unit. This is, more precisely, how Kagan understands the unitarian thesis. Kagan rejects this thesis, claiming instead that a given unit of mouse welfare is less valuable than an equivalent unit of human welfare. In economists' terms, we 'discount' mouse welfare vis-à-vis human welfare. This yields Kagan's preferred, hierarchical approach to the treatment of animals.

Why might we think that, unit for unit, human welfare is weightier than animal welfare? After all, isn't part of the point of talking of a 'unit of welfare' precisely to grant ourselves a neutral medium of moral exchange? To motivate Kagan's approach, let us start by comparing pains: assume for the moment that pain detracts from welfare equally for both mice and humans. Pain is pain, after all. For simplicity, suppose that one unit of pain deducts one unit of welfare from a life at a given moment. Now, suppose you are forced to choose between preventing one unit of pain for either a human or a mouse. If you would systematically choose to prevent the human pain, as Kagan assumes you will, that choice cries out for explanation. Either this is just an unjustified preference, or else we should reject our initial supposition that one unit of welfare is equal across species. One might try to avoid this consequence by rejecting the claim that pain detracts from the welfare of humans and mice at the same rate. But since we've assumed these to be equal units of pain, then, all else equal, it is presumably something about the respective moral status of mice and humans that makes that unit of pain less worthy of consideration in the case of mice.

On Kagan's theory, moral status is a graded affair, with humans having a higher moral status than mice. Returning to our vaccine case, this is why even the adamant vegetarian shouldn't flinch at

running vaccine trials on mice: while killing non-human animals to eat involves trading significant animal suffering for, at best, a minimal increases in human welfare, trading even equal amounts of mouse welfare for human welfare amounts to a good trade. Because mouse welfare matters less.

As noted above, Kagan contends (in chapter 7) that the same sort of quandary regarding the moral status of animals will arise for deontologists—so long as they accept that there is some number $x$ above which one is permitted to kill a single innocent to save more than $x$ lives. If the deontologist accepts such thresholds, then we can start asking what the threshold is for violating different sorts of animals' rights. Kagan is highly skeptical that the threshold for killing a mouse is going to be the same number $x$ that we have set for trading innocent human lives; rather, he thinks this number will be far lower. Again, Kagan claims, either this amounts to an unjustified preference, or else the deontologist must grant that, while mice do have a moral claim on us not to kill them, that claim is *weaker* than it is for human beings. In other words, we must accept that rights come in an ordering of strengths as opposed to being an all-or-nothing affair.

A brief methodological observation: in the course of a discussion about whether it is permissible for Tom to kill and eat a deer where this is Tom's only option for survival, Kagan rightly points out that the amount of time Tom will be lopping off the deer's life is likely to be far greater than the amount of time the flesh of that deer will serve to keep Tom alive (187). Still, Kagan takes it as a desideratum for any hierarchical theory that it vindicate the following claim: it is permissible for Tom to eat the deer in these circumstances, even if this results in an overall loss of expected welfare. A similar appeal to common sense would seem to underwrite the claim regarding the comparative disvalue of mouse as opposed to human pain. But appeals to common sense strike me as odd in the present context. Not only has Kagan himself elsewhere argued against relying on common sense claims like these in the course of doing moral theory (cf. Kagan 1989), but the claim that animals warrant no moral consideration at all was, in fact, common sense until very recently. So when exactly are we justified in appealing to common sense to support a moral theory, and when we are supposed to eschew such appeals? One will search in vain for an answer here.

This matters for the present project, because common sense judgments constitute pretty much the entirety of Kagan's evidence for his hierarchical framework. One might have expected the view to garner support from a detailed discussion of the factors which serve to ground moral status. If those factors all turn out to be gradable, one could then imagine arguing as follows: rather than trying to find an arbitrary cut-off where moral status begins, it is simpler to embrace a graded understanding of moral statuses. But while Kagan gestures towards such an argument (in chapter 5), he stops short of endorsing it—in large part because he wants to avoid taking a stand on what actually grounds moral status, beyond claiming that these grounds will have to include modal properties. That is a shame, not only because it leaves the view unsupported—aside from questionable appeals to common sense judgments—but also because it leaves us without any way of determining *how much more* human beings matter than, say, mice. Hence, it leaves us without any way of applying the framework to real-world cases, like our vaccine case.  In this respect, Kagan's book compares unfavorably to some other recent work, e.g. Budolfson & Spears (2020), which, in spite of its mere paper-length, manages to offer a detailed proposal regarding how to make such trade-offs.

In short, Kagan offers a framework for understanding how mice could matter less, morally speaking, than human beings do. But as for any explanation why they matter less, or exactly how much less they matter, that has apparently been left to the sequel.

**Bibliography**

Budolfson, Mark and Dean Spears. 2020. 'Public Policy, Consequentialism, the Environment, and Nonhuman Animals.' In *The Oxford Handbook of Consequentialism*, edited by Douglas Portmore, 592-615. Oxford: Oxford University Press.

Kagan, Shelly. 1989. *The Limits of Morality*. Oxford: Oxford University Press.

Kagan, Shelly. 2016. 'What's Wrong with Speciesism?' *Journal of Applied Philosophy* 33(1): 1-21.

Eliot Michaelson
Department of Philosophy
King's College London