

*How to Count Animals, more or less*, by Shelly Kagan. Oxford: Oxford University Press, 2019. Pp. 320

## 1. Introduction

Debates about moral status tend to focus more on necessary and sufficient conditions for moral status than on degrees of moral status. Of course, this makes sense. In a world where people still deny that sentient beings matter morally for their own sake, it is important to start by establishing that they do. Moreover, questions about degrees of moral status are difficult and sensitive. They require us to consider the possibility that individuals can have higher or lower moral status not only across species but also within species, including our own. And, no matter what view we accept, there seem to be problems. For example, are we really prepared to say that humans and ants have equal moral status? If not, are we really prepared to say that some humans can have higher moral status than others?

In this excellent book, Shelly Kagan defends a sophisticated answer to the question whether or not moral status comes in degrees. His answer is: yes and no, but mostly yes. In particular, he argues for a view that he calls ‘limited hierarchy,’ according to which (a) people have higher moral status than animals (and some animals have higher moral status than others), but (b) all people have equal moral status. At first glance, this view seems like it has no chance of working. But Kagan is a brilliant philosopher, and through a series of clever moves (ranging from how he defines terms such as ‘person’ to how he combines ideal and non-ideal theory), he makes a surprisingly strong case for his view. However, insofar as his argument rests on appeals

to intuitions about simple, idealized cases, that limits how persuasive it can be in a context where we can expect our intuitions to be deeply speciesist.

My aim in this review is to summarize the main arguments in each chapter, and then present my response. In short, I agree with nearly everything that Kagan claims in his arguments, but I disagree with the conclusions that he draws, since I disagree with the methodology that he uses to draw them. As a result, I think that this book is better at mapping the terrain than at telling us where to go on that map, though I also think that, in mapping such a complicated terrain so well, it makes an essential contribution to the literature.

## 2. Summary

In the Introduction, Kagan presents his goal for the book. Many animal ethicists accept what Kagan calls unitarianism, which holds that there is only one kind of moral status. On this view, people and animals count equally if they count at all. In contrast, common sense accepts what Kagan calls hierarchy, which holds that there are multiple kinds or degrees of moral status. On this view, people count more than animals, and some animals count more than others. Kagan has two main goals for this book. The first is to make an initial case for hierarchy over unitarianism. The second is to show how much work we need to do before we can arrive at a complete view about moral status. With that said, Kagan makes clear that his goal is not to justify our current treatment of animals: Whether we accept unitarianism or hierarchy, we should agree that our current treatment of animals is a moral atrocity that needs to end.

In Chapter 1, 'Standing,' Kagan presents a view of moral standing, that is, of what it takes to count morally for one's own sake. First, some terminology. Kagan uses 'standing' and

‘status’ differently: To have standing is to count at all, whereas to have status is to count in a particular way or to a particular degree. Kagan also uses ‘person’ and ‘human’ differently: To be a person is to be an advanced agent, whereas to be a human is to be a member of the species *Homo sapiens*. With that in mind, Kagan thinks that sentience and agency are both sufficient for moral standing. That is, he thinks that if you can feel pleasure and pain *or* set and pursue goals, then you count morally for your own sake. This means that the circle of moral concern is potentially wider than we thought: If there can be agency without sentience or vice versa, then anyone who has either counts morally for their own sake.

In Chapter 2, ‘Unitarianism,’ Kagan discusses unitarianism. On this view, everyone with moral standing has the same moral status. This might seem implausible, since, for example, if you can save a human or a mouse but not both, intuitively you should save the human all else equal. However, Kagan notes that unitarianism can make sense of the idea of different statuses too, in a way. In particular, a unitarian can hold that, if a human has more well-being at stake than a mouse, then we should favor the human for that reason. They can also hold that, if the discrepancies between humans and mice are regular and systematic, then we can describe our policy of favoring humans for this reason by saying that they have a higher moral status. So, if we want to decide between unitarianism and hierarchy, then we need to do more than show that talk of higher moral status makes sense in cases where more well-being is at stake.

In Chapter 3, ‘The argument from distribution,’ Kagan argues that, if you accept a distributive principle (such as egalitarianism, prioritarianism, sufficientarianism, or the desert-sensitive view), then you should also accept hierarchy. The reason is that unitarian distributive principles imply that we should favor animals over people, which Kagan finds implausible. Granted, a unitarian can accept this result or reject distributive principles, but Kagan finds these

responses implausible too. Alternatively, a unitarian can try to show that unitarian distributive principles do not, in fact, have this result. But Kagan argues that any such attempt will either be implausible (i.e., it will have implausible implications for people as well as for animals) or amount to hierarchy (i.e., it will imply that how much an interest counts depends in part on whose interest it is). The upshot is an argument from distribution for hierarchy.

In Chapter 4, 'Hierarchy and the value of outcomes,' Kagan discusses hierarchy. He starts by arguing for hierarchy in the context of distribution. In his view, a theory of just distribution is more plausible if people have stronger distributive claims than animals, and if some animals have stronger distributive claims than others. Kagan then argues for hierarchy in the context of the value of well-being. In his view, a theory of the good is more plausible if the interests of persons carry more weight than the interests of animals, and if the interests of some animals carry more weight than the interests of others. Kagan also discusses formulas that we can use to make the relevant adjustments (for example, we can calculate the value of your well-being by multiplying your level of well-being by your degree of moral status), and he considers and replies to objections.

In Chapter 5, 'Status,' Kagan discusses the grounds of moral status. He stresses that his view is individualist: Your moral status will depend on your individual features, not your group membership. In particular, he discusses four features that seem relevant to him. First, *well-being*: the higher your capacity for well-being, the higher your moral status. Second, *agency*: the higher your capacity for agency, the higher your moral status. Third, *potential status*: If you can have a particular moral status if the right conditions obtain in the future, then your current moral status is amplified to a degree. Fourth, *modal status*: If you could have had a particular moral status had the right conditions obtained in the past, then your current moral status is amplified to a degree.

On this view, your moral status at a given time will depend on the degree to which you have each of these features at that time.

In Chapter 6, 'Worries about hierarchy,' Kagan considers and replies to four objections to hierarchy. First, *elitism*: Hierarchy implies that some beings count more than others. Kagan replies: Every moral theory makes distinctions. The question is which are legitimate. Second, *superior beings*: Hierarchy implies that superior beings would count more than us. Kagan replies: This is acceptable, especially since these beings would still have strong duties to us. Third, *marginal cases*: Hierarchy implies that human persons count more than human non-persons. Kagan replies: This is acceptable, especially if we accept that human non-persons are 'modal persons' with amplified status, as discussed in Chapter 5. Fourth, *normal variation*: Hierarchy implies that some persons count more than others. Kagan replies: We can block this implication if we accept the theory of practical realism he discusses in Chapter 11.

In Chapter 7, 'Deontology,' Kagan discusses hierarchy in the context of deontology, according to which the rightness of an action depends on more than the value of consequences. Kagan considers two possibilities. The first is *absolutist deontology*: On this view, we have a duty to respect rights no matter how much value is at stake. The second is *moderate deontology*: On this view, we have a duty to respect rights unless a certain threshold of value is at stake. Kagan argues that both kinds of deontologist should accept hierarchy, since otherwise they will be committed to the claim that, say, killing a fish to save a human is wrong. Granted, a moderate deontologist could try to show that we meet the relevant threshold in such a case. But Kagan replies that this reply is not available to a *unitarian* moderate deontologist, since it conflicts with how we calculate thresholds in cases involving people.

In Chapter 8, 'Restricted deontology,' Kagan considers a possible response from a deontologist. We can avoid the implication that killing a fish to save a human is wrong if we accept *restricted deontology*. On this view, deontology is for people, not for animals. However, Kagan thinks that this view is implausible. The capacities that plausibly ground deontological status, such as agency and autonomy (i.e., roughly, agency plus individuality), come in degrees. And, many animals have these capacities at least to a degree. Thus, many animals plausibly have deontological status at least to a degree. Granted, we could attempt to identify a property that (a) plausibly grounds deontological status, (b) does not come in degrees, and (c) is present in people and not in animals. But Kagan replies that unless and until we find such a property (which he doubts we will), we should reject restricted deontology.

In Chapter 9, 'Hierarchical deontology,' Kagan considers another possible response from a deontologist. We can avoid the implication that killing a fish to save a human is wrong if we accept *hierarchical moderate deontology*. This view is *moderate* since it permits violating rights if a certain threshold of value is at stake. And, this view is *hierarchical* since it implies that the relevant thresholds vary with moral status. Kagan considers a variety of formulas that we might use to calculate these thresholds. He concludes that, while we need to do a lot more work to arrive at a complete theory, this kind of theory can, at least in principle, have plausible implications about a wide range of aid and rescue cases. It implies that animals have deontological status, that people have higher deontological status than animals, and that some animals have higher deontological status than others.

In Chapter 10, 'Defense,' Kagan extends this analysis to the right to defense. We normally think that, if someone attacks someone else, we are permitted to harm the attacker to defend the victim within certain limits. Kagan considers three kinds of case where this right

might apply in animal ethics: We might need to defend (a) animals against people, (b) people against animals, and (c) animals against animals. In each case, Kagan argues that hierarchical moderate deontology is more plausible than other kinds of deontology. While we again need to do a lot more work to arrive at a complete theory, this kind of theory can, at least in principle, have plausible implications about a wide range of defense cases. It implies that people and animals alike have a right to defense, but that the strength of this right varies with moral status, as well as (potentially) with other features, such as innocence.

Finally, in Chapter 11, 'Limited Hierarchy,' Kagan argues for what he calls *limited hierarchy*, which combines hierarchy with *practical realism*. According to practical realism, the moral facts are constrained by our epistemic and motivational limitations as moral agents. On this view, if it would be too demanding to think of moral status as having many levels, then we should think of moral status as having few levels instead. For example, perhaps we should have one for persons, one for near-persons, one for highly intelligent non-persons, and so on. Kagan notes that, on some moral theories, we should accept limited hierarchy as the literal truth, and on others, we should accept it as a useful fiction. Either way, we should accept it. He also notes that, if we accept limited hierarchy, we can solve the problem of normal variation, since presumably all 'normal' human persons will be at the same level.

The view that emerges for consequentialists and deontologists alike is as follows. The capacities that ground moral status come in degrees. Thus, moral status comes in degrees too. People have a higher moral status than animals, and some animals have a higher moral status than others. This is relevant both to the value of outcomes (e.g., of well-being and distributions) and to the rightness of actions (e.g., the right to aid, rescue, and defense). However, since unlimited hierarchical theory would be too cognitively and motivationally demanding, we should

accept a limited hierarchical theory instead, according to which moral status has, say, five or six levels.

### 3. Discussion

I love this book. I also somehow both agree and disagree with nearly every argument in it. On one hand, I agree with nearly everything Kagan claims about how we should interpret unitarianism and hierarchy. I also agree with nearly everything he claims about why we should reject absolutist deontology, restricted deontology, and what we might call strong hierarchy (which assigns actual, potential, and modal people much more weight than animals). In general, then, I agree that the only reasonable options are unitarianism and what we might call weak or moderate hierarchy (hereafter hierarchy, for simplicity). In exploring how these options work so thoughtfully, Kagan has done the literature a tremendous service.

On the other hand, I am skeptical of nearly every intuition that Kagan cites in support of hierarchy and against unitarianism. Throughout the book, Kagan presents us with unitarian and hierarchical options and declares that ‘few if any of us’ will agree that the unitarian option is acceptable (p. 69). In all of these cases, I disagree with Kagan about which option is acceptable. I think that everyone who has moral standing has it equally, and I am prepared to accept the consequences of this view. Granted, if a human has more well-being at stake than a mouse, then maybe we should favor the human all else equal on that basis. But we should not, in addition to favoring the human because they have more well-being at stake, *also* favor them because we assign greater moral weight to the well-being that they have.

Of course, it is not surprising that there can be a tension between ‘top down’ considerations involving considered judgments about principles such as the principle of equal consideration of interests (which I think survives the arguments that Kagan makes here) and ‘bottom up’ considerations about cases. However, not only is the methodology of this book almost entirely bottom up, but the cases that Kagan considers are almost all simple and idealized. That is, Kagan mostly focuses on cases involving simple conflicts between two individuals, such as a human and a fish. Additionally, when Kagan discusses these cases, he mostly abstracts away from the real-world conditions that make human and nonhuman lives so different.

I want to be clear that I think that this kind of thought experiment can be useful. Considering our intuitions about simple, idealized cases can be a good way of clarifying our concepts and the relevance of particular variables (especially since, as Kagan so effectively demonstrates, even the simplest cases have hidden complexity). However, this should only be part of what we do when we construct moral theories. And when we are considering a topic like animal ethics, when our intuitions are so heavily influenced by speciesism and other such biases, there is a risk that focusing narrowly on simple, idealized thought experiments, as Kagan does here, will anchor us to an unacceptably conservative and speciesist moral theory.

For example, consider just some of the potentially misleading considerations that our intuitions about these cases might be drawing from.

First, we intuitively underestimate the capacities of nonhuman animals for a variety of reasons. We tend to perceive happiness and autonomy more easily in humans than in other animals. Moreover, insofar as there are limits on how happy or autonomous an animal can be, we tend to attribute these limits to internal causes, i.e., facts about the animal, rather than to external causes, i.e., facts about the conditions in which the animal is living. If we were to correct for

these tendencies, then we would likely see less of a divide between humans and other animals than we currently do.

Second, we intuitively underestimate the moral status of nonhuman animals for a variety of reasons as well. Even if we were to correct for our mistaken empirical assumptions, we would still be likely to see animals as having a lower moral status than they do because we have a bias in favor of members of our own species, and because we know that if we were to assign a higher moral status to animals, then we would have to change nearly everything about how we live our lives and structure our societies.

Of course, we have these tendencies in the case of other humans too. Insofar as other people are different from us (and especially insofar as other people are less privileged than us), it can be easy to underestimate their capacities and moral status. It is therefore not surprising that we would do the same with animals. And, since differences across species are even greater than differences within our species, it is also not surprising that we might underestimate many animals even more than we underestimate many humans.

I realize that I am mostly stating the obvious here. But I think that stating the obvious is important, because it shows how deeply unreliable our intuitions are in this context. Insofar as we underestimate animals empirically, our intuitions in support of hierarchy are suspect, because we will be intuitively overestimating the difference between human and nonhuman animals, and so we will be intuitively overestimating the plausibility of hierarchy as a partial explanation of that perceived difference. And of course, insofar as we underestimate animals normatively, our intuitions in support of hierarchy are once again suspect, this time for direct and obvious reasons.

Kagan might be happy to concede this point, and to accommodate it by noting that, if we expect to intuitively underestimate animals, then we can 'round up' our intuitions accordingly.

But this will not be enough to address the issue, since our intuitions about how much to ‘round up’ our intuitions about animals will themselves be shaped by our intuitions about animals. For this reason, I think that we need to do more than correct our intuitions in this context. We also need to limit the role that our intuitions play in this context, by allowing ‘top down’ considerations to do more of the work.

Kagan might be happy to concede this point as well. He claims throughout this book that we need to consider many more issues before we can decide which theory is best all things considered. I appreciate this qualification. However, I still worry that his book will make the hierarchical theory appear to be more plausible than it is by focusing our attention on intuitions that should, at best, be playing a minor role in shaping our considered judgments about which theory to accept. So, even if this book is meant to be only one small part of a much larger discussion, I still worry about how it might shape the trajectory of the conversation.

To be clear, when I claim that a simple, idealized, bottom up methodology in the context of speciesism can bias our moral thinking, I am not only claiming that it can lead us to overvalue hierarchical views with steep discount rates. I am also claiming that it can lead us to overvalue hierarchical views as a general matter. I believe that if we considered principles as well as cases, from ideal as well as non-ideal perspectives, while correcting for bias in our descriptive as well as our normative intuitions, we would find that the perceived gap in intrinsic moral status between, say, a human and a fish not only narrows but closes entirely.

For these reasons, my own preference is to read this book not primarily as an argument for the hierarchical theory over the unitarian theory, but rather primarily as an exploration of what these theories imply on different interpretations, with running commentary about what Kagan happens to find intuitively plausible. I am also choosing to interpret this running

commentary not primarily as evidence of which theory is correct, but rather primarily as evidence about what we think about animals, which is helpful to keep in mind so that we can identify and correct for possible biases in our thinking about animal minds and ethics.

I will now close by briefly making three other, related points.

First, there is an ambiguity in our current discourse about moral status that, I worry, might affect how people assess unitarianism and hierarchy. Some people see moral status in fully intrinsic terms, as a matter of what moral claims we have in virtue of our capacities. Other people see moral status in at least partly extrinsic terms, as a matter of what moral claims we have in virtue of our capacities, contexts, and/or relationships. Why does this matter? It matters because, insofar as we conflate these interpretations of moral status, there is a risk that we will see humans as having a higher intrinsic moral status than other animals at least in part for extrinsic reasons. For example, we might think, rightly or wrongly, that we have stronger duties to humans than to other animals at least partly on grounds involving cluelessness (we know more about how to help humans), demandingness (we can help humans more easily), special relationships (we have reciprocal relationships with other humans), and so on. If so, we might also have the intuition that humans have a higher moral status than other animals, all else equal, without clearly distinguishing the intrinsic and extrinsic contributors to this intuition.

Second, throughout the book Kagan seems to assume that, if we accept unitarianism, then we should extend our views about the moral status of humans to nonhumans. Thus, for example, if we accept a certain formula for determining when we can harm humans in cases of aid, rescue, or defense, then we should accept roughly that formula for nonhumans as well. However, I think that this is a mistake. If we accept unitarianism, then we are accepting that our final theory of moral status will be the same for anyone with moral standing, human and nonhuman alike.

However, as Kagan repeatedly emphasizes, taking seriously the moral status of nonhumans is a game-changer. Yes, we should think about what makes sense in the case of humans, and how that would extend to nonhumans. But we should also think about what makes sense in the case of nonhumans, and how that would extend to humans. When we do, we might realize that the truth is somewhere in the middle and, as a result, is much different than what seemed intuitively plausible about either case on its own. (And of course, we might change our view yet again when we start taking other metaethical and normative considerations seriously as well.)

Finally, and perhaps most importantly, even if Kagan were right that we should accept hierarchy, it would still be the case that we should prioritize nonhumans over humans in many cases in practice. For consequentialists (and some non-consequentialists), this is because nonhuman suffering is so massive, neglected, and tractable. And, for non-consequentialists (and some consequentialists), this is because we are responsible for a lot of nonhuman suffering, we benefit from the activities that cause a lot of this suffering, and we can address much of this suffering without substantially reducing our own quality of life. Thus, any reasonable hierarchical view should imply that nonhuman suffering is a moral priority for consequentialists and non-consequentialists alike. Granted, a reasonable hierarchical view might still imply that we should prioritize humans in many cases in practice. But it will also have the opposite implications in many cases in practice, especially when we consider our impacts on entire nonhuman populations. The world is a terrible place for animals, in very large part because of us, and so we have a responsibility to prioritize them whether or not we think that we have a higher intrinsic moral status than they do.\*

JEFF SEBO  
New York University

[jeffsebo@nyu.edu](mailto:jeffsebo@nyu.edu)

\* I am grateful to Dale Jamieson, Shelly Kagan, and the Editors of *Mind* for helpful feedback on the penultimate draft.