



ARTICLES

Intuition and Kagan's Hierarchicalism

Xihe Ouyang 

Department of Philosophy, University of Calgary, Calgary, AB, Canada
Email: xihe.ouyang@ucalgary.ca

Abstract

According to Shelly Kagan, the moral status of an individual is determined by the extent to which the individual has (has now, might/will have, or could have had) certain psychological capacities. Roughly speaking, the greater one's relevant psychological capacities, the higher their moral status. In this paper, I offer a twofold critique of Kagan's hierarchicalism. On the one hand, I argue against the primary argument in favor of Kagan's view (the argument from distribution) by challenging the key intuition on which the argument relies, thereby reducing the appeal of Kagan's position. On the other hand, using Kagan's general methodology, I argue that a good reason to reject Kagan's account of moral status is that he fails to explain away the counterintuitive result of his theory in the case of normal variation.

Keywords: moral status; Shelly Kagan; hierarchicalism; the argument from distribution; practical realism

Introduction

A being has *moral standing* if and only if it counts, morally speaking, in its own right and for its own sake. The notion of moral standing is crucial from the perspective of ethics because only beings that have moral standing can be wronged. Some common-sense examples may help to elucidate this point. We believe, for instance, that ordinary humans can be wronged, and that they morally count in their own right. By contrast, we do not tend to think that by destroying a rock (say) the rock itself is wronged. Although we might think that the rock's owner is wronged, we would not recognize a rock has having moral standing.

Having moral standing means that the relevant being morally counts. Nevertheless, it is frequently a matter of contention whether all individuals with moral standing count equally. In this context, people often appeal to the concept of moral status.¹ Moral status is a normative profile that governs how other moral agents are morally required to treat a being that has that moral status.² Having moral standing entails having a moral status.

¹This refers to default moral status, that is, the moral status we assign to a being before we consider its (past and possible future) morally praiseworthy or blameworthy behaviors.

²Some philosophers, such as Peter Vallentyne, believe that moral status can be bifurcated into two categories: moral agent status and moral patient status. The former is a particular kind (or level) of moral standing that governs how a moral agent with that kind (or level) of moral standing is morally required to act. Moral patient status is a particular kind (or level) of moral standing that governs how other

Typically, compared with beings with lower moral status, a being with higher moral status has more or stronger relevant rights and/or their interests carry greater weight in moral calculus (several examples will be introduced throughout the paper).

On the issue of moral status, Shelly Kagan (2016, 2019) defends a version of what is called 'hierarchicalism.' He believes that there is more than one moral status, and that the moral status of some beings is higher than that of others. According to him, the moral status of an individual is determined by the extent to which the individual has (has now, might/will have, or could have had) certain psychological capacities. Roughly speaking, the greater one's relevant psychological capacities, the higher their moral status.

In this paper, I will argue against Kagan's view. My project contains two main parts. On the one hand, I argue against the primary argument in favor of Kagan's view (the argument from distribution) by challenging the key intuition on which the argument relies, thereby reducing the appeal of Kagan's stance. On the other hand, employing Kagan's general methodology, I argue that a good reason for rejecting his view of moral status is that he fails to explain away the counterintuitive result of his theory in the case of normal variation.

1. Kagan's hierarchicalism

Kagan believes that one's moral status is entirely determined by a set of intrinsic properties that one has (has now, might/will have, or could have had), rather than one's membership in certain groups or one's relationships with other beings. More specifically, he argues that an individual has moral standing if that individual has the capacity for sentience (which is understood in a narrow sense as the ability to experience pleasure and pain) or agency (which is understood as having preferences, even in a very minimal sense, for how one's life should go).³ The capacity to experience pain or pleasure is relevant to the moral status of a being insofar as it implies the possession of well-being (in a narrow sense) that can be increased or reduced. Agency, on the other hand, is relevant to the moral status of a being because its possession suggests that the relevant being has a will that merits respect. The possession of either of the two capacities is sufficient for ensuring that the relevant being has a level of moral status.

According to Kagan, actually possessing the relevant capacities at a certain time is necessary for a being to have moral standing and a moral status at that time. If a being does not have any relevant capacities at time T, that being has no moral standing and no moral status at time T. Nevertheless, as briefly mentioned above, Kagan argues that actual capacities are not the only things that can affect a being's moral status. *Potential capacities* and *modal capacities* (as defined below) can raise one's moral status to a certain degree, although their possession is not sufficient for a being to have a moral status.

Potential capacities are the capacities that a being does not currently possess, but that it has the potential to develop in the future. Modal capacities are the capacities that, in the past, the individual had the potential to have now, but which the individual currently lacks. Kagan holds that potential and modal capacities can influence a being's moral status, but only if the actual capacities of that being are sufficient to already grant them a

moral agents are morally required to treat a being that has that kind (or level) of moral standing. In this paper, I will focus on moral patient status, and the term "moral status" will denote moral patient status.

³Given that Kagan focuses on the cognitive parts of psychological capacities rather than the emotional parts, I will do the same in this paper.

certain moral status. Two examples might be helpful at this point. First, regarding potential capacities, consider the case of a sentient baby. A sentient baby has a moral status because of their sentience. Moreover, although the baby does not have the capacity to read novels now, they have the potential to possess that ability in the future. According to Kagan, this potential capacity can increase the baby's moral status. Second, regarding modal capacities, consider the case of Tom, who is now cognitively impaired as the result of a car accident 10 years ago. It *could have been* the case that Tom avoided that car accident 10 years ago and remained cognitively unimpaired. In this case, according to Kagan, we can say that Tom has a modal capacity to be a cognitively unimpaired adult, and this modal capacity may raise his moral status.⁴

Another hallmark of Kagan's view is his endorsement of *hierarchicalism*. Roughly speaking, Kagan thinks that, all else being equal, the degree to which a being possesses morally relevant psychological capacities determines how much it matters morally. The possession of more advanced psychological capacities typically implies a higher moral status, whereas an equal level of the same capacities entails an equal moral status. This view is a version of hierarchicalism insofar as it proposes a hierarchy of moral statuses with some beings' moral status being higher than that of others.

To better understand hierarchicalism, we might contrast it with *unitarianism*. Unitarianism suggests that every being with moral standing has the same moral status. One version of this is the equal consideration of interests found in Peter Singer's work (1990, 2009). According to Singer, if two beings have exactly similar interests at stake, these interests have the same weight in the relevant moral calculus. For example, given that both an ordinary pig and an ordinary human are sentient, Singer suggests that if they have exactly similar interests at stake, then their interests should weigh the same, regardless of who has stronger psychological capacities. Conversely, according to hierarchicalism, it is possible for two beings to have different and hierarchically ordered moral statuses. In this case, then, the interests of the being with a higher moral status will carry greater weight than the exact similar interests of the being with lower moral status.⁵ This implies that the goodness of pleasure or the badness of pain depends on who feels it. For example, the pain endured by a person weighs more in the relevant moral calculus than the exact similar pain felt by a pig if the former has a higher moral status.

In summary, Kagan's view is an individualistic, capacity-based hierarchicalism. However, despite the view's intuitive appeal, I will argue against it. I begin by challenging the strongest argument Kagan provides in favor of his view, namely, the argument from distribution.

2. The argument from distribution

To do justice to Kagan's hierarchicalism, it is necessary to consider his main argument. The strongest argument he provides to show the superiority of his view over unitarianism is the argument from distribution.⁶ The argument begins by noting that if animals

⁴Kagan labels the view that modal capacities are relevant grounds of moral status "modal personism." See Kagan (2016).

⁵Here, I assume that moral goodness is affected by status. A similar logic also applies to moral permissibility.

⁶This argument provided by Kagan is inspired by McMahan (1996) and Vallentyne (2005). Kagan, of course, also gives other arguments in favor of his view. For example, building on Kazez (2010), Kagan

possess the same moral status as cognitively unimpaired human adults (as unitarianism maintains), our distributive principles should also extend to them equally. However, if this is the case, we may be required to put most of our resources and efforts into improving the well-being of animals, as most animals have (much) lower levels of well-being than humans, and thus may have priority over those who are more well-off. Such an outcome, Kagan contends, is counterintuitive and overdemanding, rendering unitarianism implausible.

To elucidate the argument, let us consider a popular version of egalitarianism that advocates for minimizing the inequality between individuals' well-being. If animals were to have the same moral status as unimpaired human adults, the egalitarian principle would be applied to them equally. This implies that we ought to minimize the inequality of wellbeing between animals and humans. Now, given that, as noted above, the level of well-being that most animals have is significantly lower than that of most humans, we would then be required to make a remarkably increased effort to promote the well-being of animals.⁷ This includes (but is not limited to) redirecting a significant portion of the resources that are under our control toward animals. This result, Kagan suggests, is absurd. Our intuition suggests that *the requirement that we should put most of our resources and efforts into improving animals' well-being is overdemanding and thus implausible*. (Call this intuition **K-intuition**.)

The argument from distribution does not apply only to egalitarianism. As Kagan remarks, a similar absurd result will be obtained if we replace egalitarianism with any of the three other most plausible theories of distributive justice – prioritarianism, sufficientarianism, and desertism – as long as unitarianism is true. Given that we have good reasons to believe that at least one of those distributive principles is plausible and given that those principles combined with unitarianism imply the counterintuitive result (K-intuition), we should conclude that unitarianism is implausible. On the other hand, Kagan argues that his hierarchical view remains unaffected by this argument. Given that his view suggests that animals have a lower moral status than us, we are not required to put most of our resources and efforts into improving their well-being, no matter what distributive principles we endorse.

Some objections against the argument from distribution have already been put forward in the literature.⁸ However, rather than discussing or developing any of the existing criticisms, I offer a new objection that centers on the legitimacy of K-intuition. In particular, I will argue that there are compelling reasons for believing that K-intuition is biased and unjustified. Consequently, I will maintain that K-intuition provides no good reason to believe that the result brought about by the combination of unitarianism and

provides a case to show that intuitively, the well-being of a being with higher relevant capacities matters more than the exactly similar amount of well-being of a being with lower relevant capacities. Moreover, he thinks that by introducing modal capacities, his view sidesteps certain criticisms commonly directed at the traditional version of hierarchicalism (e.g., the problem of marginal cases). I believe that among all the arguments in favor of Kagan's view, the argument from distribution is the strongest one.

⁷Note that this requirement can hold even when the relevant animals are only minimally sentient.

⁸Sebo (2021) questions the general methodology of the argument and claims that intuition should not play such an important role in moral reasoning. Brouwer and van der Deijl (2020) argue that compared with humans' well-being, animals' wellbeing levels are not as low as we might imagine. Moreover, they also argue that we lack the capacity to effectively improve the level of wellbeing of many animals, and this practical limitation implies that we do not need to significantly reform our society to benefit animals. Personally, I believe that these three objections fail to undermine the plausibility of the argument from distribution. However, given the limitations of space, I do not argue against them in detail.

distributive principles – namely, that we are morally required to put most of our resources and efforts into improving animals’ well-being – is objectively wrong or implausible. Thus, I will claim that the argument from distribution does not give us a good reason to reject unitarianism.

3. K-intuition and the bias of the powerful

Although many may find K-intuition appealing, I believe that there is a good reason to suspect that this intuition is biased. It is highly likely that the reason we believe the requirement to put most of our resources and efforts into improving animals’ well-being would be overdemanding and thus implausible is that we are influenced by the general bias of the powerful.

To illustrate the bias of the powerful, let us first consider cases within human society. Typically, if the powerful are asked to give away most of what they have to meet a relevant moral requirement, they will find it counterintuitive. Take the case of ancient Chinese emperors. If one were to ask an emperor from ancient China to give away most of what they had to ordinary people, they would have found the request absurd. Even for beneficent emperors, the idea of redistributing their wealth to achieve societal equality or prioritize the needs of the powerless would have seemed outlandish. Presumably, we would get a similar reaction if we were to ask the same of other highly powerful individuals who lived before the Enlightenment. Ancient slave owners and Pharaohs are both cases in point. Obviously, given our current moral knowledge, we know that this kind of reaction or intuition is biased and unjustified. Similarly, it is unsurprising that when the human community (which is extremely powerful in comparison to animals) is asked to give away most of what they have to animals to meet a relevant moral requirement, they tend to find it counterintuitive.

The bias of the powerful has not gone unnoticed among philosophers, especially pragmatists, who provide us with theoretical explanations of it. By considering these explanations we can judge whether the same mechanism that underlies the bias within human society also applies to that between humans and animals. Following John Dewey (1932), Elizabeth Anderson (2014, 2021) points out that the bias of the powerful, as a kind of self-serving bias, is ubiquitous. Moreover, generally speaking, the more powerful one is, the more susceptible they are to such biases. Anderson also remarks that convincing empirical evidence suggests, “Standing in a position of superior power over others tends to bias the moral sentiments of the powerful, in at least three ways: it reduces their compassion, activates their arrogance, and leads them to objectify subordinates” (Anderson 2021: 88). One general upshot of this bias is that the powerful often believe that the existing social structure that favors their interest is also the morally right one. In other words, they tend to conflate what is morally right with what is beneficial to them.

Let us now compare the relationship between the powerful and the powerless within human society (the human-human case) with the relationship between powerful humans and powerless animals (the human-animal case). There are several reasons to believe that the bias of the powerful is not only operative in the human-animal case, but more severe than it is in the human-human case.

First, the power difference is evidently more pronounced in the human-animal case than it is in the human-human case. As previously noted, typically, the more powerful one is, the stronger their self-serving bias will be. Thus, it is highly likely that the powerful in the human-animal case holds a stronger bias than the powerful in the human-

human case. Second, the objectification of animals is a widespread phenomenon in the Western world. Moreover, the degree to which humans objectify animals is arguably more intense than the degree to which powerful humans objectify powerless humans. Historical evidence underscores this point: powerful men have historically sought to objectify women and slaves by likening them to animals. Third, compared with the human-human case, in the human-animal case, it is harder for the powerful to notice the needs, feelings, interests, and rights of the powerless: all humans share similar biological structures and can relate to each other's subjective experiences, but the same cannot be said for humans and animals. Moreover, the absence of a shared language with animals further complicates our understanding of their needs and emotions, especially for those without a background in ethology. Fourth, the challenges in recognizing the needs and emotions of animals make it more difficult for humans to empathize with them. Given the pivotal role of empathy in shaping moral intuitions, this lack of empathy can lead to biased intuitions.⁹

What is more, the bias in the human-animal case is not only stronger than in the human-human case but also harder to be corrected. In the human-human case, the powerless can usually assert their rights and needs to correct the bias of the powerful, making the latter acknowledge them. This can be achieved through various forms of resistance, ranging from petitions and lobbying to more assertive measures such as boycotts or even revolutions (Anderson 2014). In this way, the powerful are pushed to understand that the powerless are not mere objects, but their moral equals, and that the existing structure that systematically favors the powerful is not a morally right one.

Unlike powerless humans, by contrast, animals cannot organize social movements or revolutions to force the powerful to acknowledge and respect their feelings, needs, rights, and interests.¹⁰ Moreover, I believe that Anderson rightly points out, “[i]ndividuals learn the difference between what they want and what is right through experiences in which others hold them to account for wrongdoing through practices such as blame and punishment” (Anderson 2021: 88). And yet, in the human-animal case, animals are hardly in a position to hold humans accountable through blame and punishment. Although there might be humans who champion animal rights and initiate movements on their behalf, the endeavor is undeniably more challenging.

To summarize, there are good reasons to believe that the bias of the powerful not only exists in the human-animal case, but that it is even stronger and more pernicious than in the human-human case.¹¹ This general bias, which unjustifiably favors humans over animals, is highly likely to be responsible for the formation of K-intuition. If this is

⁹Harry. R. Lloyd also mentions this when he argues against the value of the wellbeing argument. See Lloyd (2021).

¹⁰Hume suggests something similar. See Beauchamp, Tom L. (ed.) (1998). *An Enquiry Concerning the Principles of Morals*, M 3.18, SBN 190-1 and M 3.19, SBN 191.

¹¹One may wonder how the consideration of self-interest fits into this picture. It could be argued that the bias of the powerful is partly caused by considerations of self-interest. The unequal power relationship amplifies the bias of self-interest in the powerful as the powerless lack the capacity to prompt the powerful to consider the interests of the powerless. On the other hand, the powerless are less likely to conflate their self-interest with moral norms because they are forced to consider other people's interests all the time. In this way, one may even argue that the consideration of self-interest plays an important role in forming K-intuition by helping to generate the bias of the powerful. Alternatively, one may hold the view that the bias of the powerful is just one specific and strong form of the bias of self-interest. On this issue, there is certainly more that can be said, but given the scope of the paper, I cannot explore it further here. It is crucial to note, however, that these suggestions do not undermine the argument I present in this section. All I

the case, then K-intuition is highly likely to be biased and unjustified. Therefore, such an intuition must be discarded as a reason for believing that the combination of unitarianism and distributive principles will lead to an objectively implausible result. Once the force of K-intuition is significantly undermined, the argument from distribution becomes much less plausible.¹²

4. The problem of normal variation

Having undermined the key intuition to which the argument from distribution appeals, in the rest of the paper, I continue to challenge Kagan's view by presenting cases in which people's intuitions, rather than supporting his theory, strongly contradict it.

As already mentioned, Kagan's argument from distribution is greatly dependent on people's intuitions. Although, in some places, Kagan seems to claim that the argument from distribution is so decisive that we should reject unitarianism because of it (Kagan 2019: 76), I believe a more plausible and charitable reading of Kagan's argument is as follows: given the plausibility of our distributive principles and the fact that K-intuition is widely shared, we have a compelling *pro tanto* reason to abandon unitarianism. On the one hand, this reading suggests that whether we should abandon unitarianism at the end of the day also depends on other factors in our reflective equilibrium, which, methodologically speaking, is more plausible. On the other hand, it preserves the key methodological point on which Kagan relies; namely, if a view contradicts a widely shared intuition, then we will have a good *pro tanto* reason to reject that view. Nevertheless, one may still question the methodological principle behind this charitable reading (see Sebo 2021). Personally, I am sympathetic to this methodological principle, but a problem for Kagan is that by employing the same methodological principle he uses, we can form arguments against his view. In other words, we can form cases in which people's intuitions strongly contradict Kagan's hierarchicalism. One such challenge to Kagan's view is the problem of normal variation, to which we now turn.

The problem of normal variation can be described as follows: among unimpaired human adults, there is variation in the extent to which individuals possess the relevant cognitive capacities (including current, potential, and modal capacities). There is, in other words, a variation among unimpaired adults in the capacities that are relevant to their moral status. In light of this, Kagan's hierarchicalism would imply that there are differences in moral status even among cognitively unimpaired adults. For example, all else being equal, if one has a poorer memory than others, Kagan's hierarchicalism seems to suggest that that person counts less than others in the relevant aspect (Kagan 2019: 165). Arguably, most people will find this implication unacceptable

need to show is that there are good reasons to believe that the strong bias of the powerful exists in the human-animal case and that it is highly likely to be responsible for the formation of K-intuition.

¹²My argument against K-intuition does not automatically imply that the opposite of K-intuition – that we should redistribute massively to help animals (the redistributive claim) – is correct. (For example, one may think both K-intuition and the redistributive claim are problematic and decide to stay agnostic on this issue.) Although I personally believe that we should redistribute, I admit that to seriously defend such a claim requires further efforts, which I do not provide in this paper. Nevertheless, the argument put forward in this section does not depend on a commitment to the redistributive claim. Instead, all I need to show is that K-intuition is not as plausible as it seems. To make the argument from distribution work, one needs to show that K-intuition is (much) more plausible than the redistributive claim, and as I have tried to show, this is far from obvious.

insofar as it is generally thought that all unimpaired adults possess equal moral status.¹³ Thus, we have a compelling pro tanto reason to reject Kagan's view.

Kagan is aware of this problem, and he attempts to address it by introducing two solutions, namely, practical concern and practical realism. However, as I will argue, neither of the two solutions is tenable, leaving the problem of normal variation unsolved. This problem, then, continues to be a severe challenge to Kagan's view.

4.1 Practical concern and practical realism

Kagan suggests that his hierarchicalism offers two possible ways to solve the problem of normal variation. The first is *practical concern*, which suggests that the relevant differences in moral status among cognitively unimpaired adults are so small that they can be disregarded for *practical purposes*. However, Kagan himself admits that most people would find it unacceptable to embrace "differences in moral status, no matter how small, among ordinary nonimpaired, adult human beings. ... even when it has little or no practical import" (Kagan 2019:166). That is, most people believe, *as a matter of fact*, that there are no differences in moral status among unimpaired adults and reject the thought that we should merely act as if there were no such differences. Given this widespread intuition, we must dismiss practical concern as a plausible solution.

Kagan's second strategy to address the problem of normal variation, *practical realism*, suggests that when we *generate* moral rules or principles, we should be realistic about the epistemic and motivational capacities we have. According to this strategy, given our epistemic and motivational limitations, it is unrealistic to expect us to discern and work with a highly complex and fine-grained system of moral statuses. Instead, practical realism suggests that we adopt a relatively simple system of moral statuses, one that we can identify and reliably follow. This approach advocates for a hierarchy with only a limited number of moral statuses as opposed to a countless number of statuses.

Importantly, according to practical realism, this coarse-grained hierarchy provides us with *literal moral truths* rather than a mere heuristic or useful fiction. To understand what "literal moral truth" means, consider an analogy between practical realism and *realistic rule consequentialism*. Realistic rule consequentialism tries to factor in human limitations when formulating moral rules. If the rules are too complicated or demanding, they are unlikely to yield the best overall outcome, because people cannot follow them. Thus, realistic rule consequentialism advocates the adoption of a set of optimal rules with limited complexity. Given that these rules are compatible with human epistemic and motivational limitations, they allow us to achieve the real best outcome. Crucially, according to realistic rule consequentialism, these rules are literally true rules in our moral system, not a heuristic of some further rules. In a similar way, practical realism holds that our psychological constraints only allow us to recognize a relatively limited number of moral status levels, and thus the corresponding limited hierarchy represents a *literal moral truth*.

This approach can be illustrated in a more straightforward way by comparing two charts (Figures 1 and 2). In both charts, the *x*-axis stands for the level of relevant psychological capacities, and the *y*-axis represents the corresponding moral status. However, the chart in Figure 1 relates the variation in relevant psychological capacities and the corresponding variation in moral status through a step function, whereas the chart in Figure 2 relates them through a linear growth function. Hierarchicalism

¹³Personally, I also share the intuition that all unimpaired adults possess equal moral status. Moreover, even Kagan admits that this intuition is shared by nearly everyone (Kagan 2019: 166).




Figure 1. Hierarchicalism with practical realism.

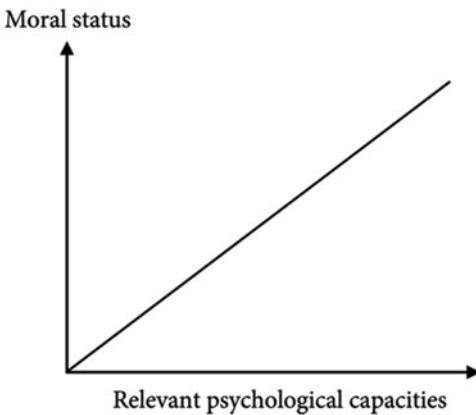


Figure 2. Hierarchicalism without practical realism.

without practical realism would favor a linear growth function, where every slight increase in relevant psychological capacities corresponds to an increase in moral status. In contrast, practical realism rejects a linear growth function because our cognitive capacities are not equipped to identify countless numbers of moral status levels and because it would be practically impossible for us to follow and apply such a nuanced system in real-world moral decision-making. According to practical realism, the correct relation between the variation in relevant psychological capacities and the corresponding variation in moral status is the one captured by a step function, which implies that for variations over vast ranges of capacities, the moral status remains constant, jumping to a higher level only after a significant increase in capacity.

Practical realism suggests a step function that we can handle. As shown in [Figure 1](#), “many differences in capacity will make no difference to status at all, for there will only be a small number of possible levels at which any given individual can be placed” (Kagan 2019: 289). It is then very likely that the variations in capacities among unimpaired adults are too small to generate differences in moral status. Thus, all unimpaired adults would be assigned the same moral status, which, according to practical realism, would be literally true. In this way, the problem of normal variation is made to disappear.

Before assessing the merits of practical realism, it is useful to bring out what distinguishes it from practical concern. Practical concern suggests that the literal moral truth is that there is a countless number of moral statuses, but that we act *as if* there were only a limited number of them, and as if all unimpaired adults have the same moral status. In contrast, practical realism suggests that it is *literally true* that there is only a limited number of moral statuses, and it is also *literally true* that all unimpaired adults have the same moral status. Kagan is right about the fact that most people, including me, would not accept practical concern. However, I believe that practical realism, too, fails to solve the problem of normal variation.

4.2 Why practical realism is problematic

4.2.1. Moral status relativism: species

According to practical realism, the literally true system of moral statuses depends partly on the psychological capacities of the group of beings in a given society at a given time. So, facts about moral status are moral-agents-mind-dependent. In other words, certain features of some moral agents' minds (e.g., their epistemic capacities) partly determine how many levels of moral status there are in the world.¹⁴ This is puzzling. Why should certain moral agents, who are far from cognitively perfect (e.g., humans), set the standard for how many levels of moral status there are in the world? Besides this puzzlement, this feature of practical realism leads to a more fatal problem.

Practical realism implies that different species with different psychological capacities may have different literally true moral status systems. Other species whose psychological capacities are different from ours can develop their own literally true moral status systems, which are likely to be different from ours. As a result, an individual may possess one moral status within one system and a different status in another. Given that both systems convey literal truth, we are confronted with a form of moral status relativism.

The problem can be corroborated by the following thought experiment. Imagine there is an alien species whose normal members are cognitively much superior to humans. One day, an alien belonging to this species meets a human, Luka. They have the following conversation.

Luka says: "Peter and I are unimpaired humans. Though there are some morally relevant cognitive differences between us, these differences are too small for us to reliably discern and handle them. Thus, according to practical realism, Peter and I have the same moral status" (see Figure 3).

The alien replies: "I accept practical realism, but I believe that Peter has a higher moral status than you. Unlike your primitive brain, my brain and my fellow species members can recognize and deal with over 100000 levels of moral status; and I detect that, other things being equal, Peter has slightly better cognitive capacities than you. According to my hierarchy, you and Peter belong to different levels of moral status, his status being higher. If your moral status can be represented by 1 unit, then his moral status is 1.1 units (see Figure 4). The reason why you assign Peter 1 unit of moral status is that your moral status system is not fine-grained enough."

¹⁴Note that in early chapters, Kagan asserted that moral status is based only on the relevant moral patient's intrinsic properties. However, embracing practical realism makes moral status partly extrinsic, since it partly depends on (other) relevant moral agents' psychological capacities. This may be another problem for Kagan and his practical realism.

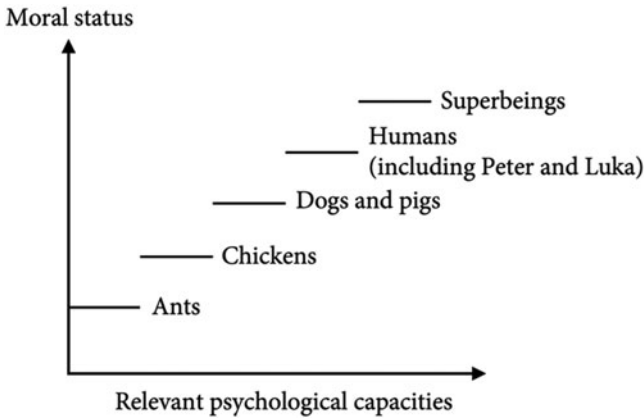


Figure 3. Luka's moral status system.

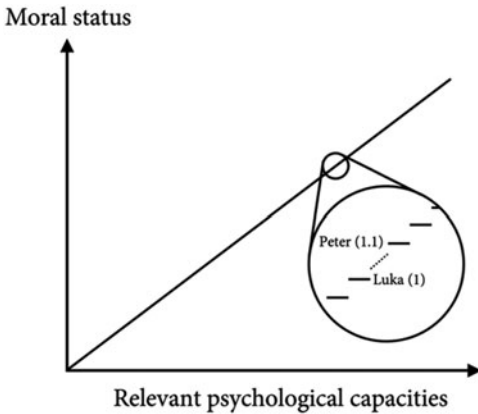


Figure 4. Alien's moral status system.

How should we interpret this thought experiment?¹⁵ If Kagan's hierarchicalism and practical realism are true, then both Luka and the alien are correct. Luka has one moral status system while the alien adheres to another. Yet, despite the significant discrepancies between the two systems, both are literally true. To avoid logical contradiction, we would need to resort to moral status relativism. That is, different species with different capacities will have different literally true moral status systems, potentially leading to different judgments about the same individual's moral status.¹⁶ In the given case, for

¹⁵It seems to me that the most plausible thing we can say is that both are wrong. Facts about moral status are objective and mind-independent. Although the alien could discern a difference between Luka's and Peter's moral status and the difference between chickens' and pigs' moral status, it is likely that there are even more subtle differences that they cannot discern. They are merely closer to the objectively true moral status system (assuming hierarchicalism is correct).

¹⁶One may argue that even different individuals belonging to the same species but endowed with different cognitive capacities may have different moral status systems. However, two remarks are in order. First,

the alien, the moral status of Peter is 1.1, whereas, for Luka, Peter has a moral status of 1. Moral status relativism implies that both judgments are true. However, it is counter-intuitive to suggest that, as a matter of fact, Peter has two different moral statuses at the same time (and he would have even more statuses if we considered other species of moral agents).

If one still doubts why moral status relativism is problematic, consider the following additional case. For the sake of the argument, let us suppose that utilitarianism is true.¹⁷ Imagine a scenario in which Peter and Luka are in danger, and in which only one of them can be saved by a doctor. Suppose the unweighted value of Peter's continued life is 1 and the unweighted value of Luka's continued life is 1.05. Now, let us take their moral status into account and calculate the weighted value of their continued life (the weighted value of their continued life = the unweighted value of their continued life \times their moral status). Recall that according to the alien's moral status system, the moral status of Peter is 1.1, whereas Luka's moral status is 1. If the alien is correct, the weighted value of Peter's continued life would be 1.1 (1×1.1), and the weighted value of Luka's continued life would be 1.05 (1.05×1). Other things being equal, this result indicates that the doctor should save Peter. However, in Luka's system, the moral statuses of both Peter and Luka are 1. If Luka's system is correct, the weighted value of Peter's continued life would be 1 (1×1), and the weighted value of Luka's continued life would still be 1.05 (1.05×1). Other things being equal, this result indicates that the doctor should save Luka. Recall that according to practical realism, both the alien and Luka's moral status systems are literally true. However, the alien's system says we should save Peter, but Luka's system recommends his own salvation. It strikes me as implausible to suggest that the opposed recommendations are literally right at the same time. Thus, moral status relativism is implausible.¹⁸ If moral status relativism is implausible, then practical realism, and consequently Kagan's hierarchicalism, are also called into question.

different normal individuals who are members of the same species are likely to have similar cognitive capacities, so according to practical realism, they probably have the same moral status system. Second, it is reasonable to suggest that in a one-species-community, members should reach an agreement on the moral status system that can and should be followed by its members. Thus, the mere consideration of different individuals within a single species need not entail moral status relativism (within the corresponding species). Moreover, even if it were the case that different individuals within the same species had different moral status systems, this agent-relative proposal would only make things worse. If the moral status system is agent-relative, then we will encounter implausible moral status relativism even on the individual level.

¹⁷The general idea of the argument still holds when other normative theories are correct, providing we modify the details of the argument accordingly.

¹⁸Of course, I acknowledge that not everybody finds this relativism implausible. For example, an anonymous reviewer of this journal claims that if one accepts rule consequentialism, and one accepts that the best set of rules for one community (e.g., species) need not be the best for a different community, then one may hold that the relativism which follows is not absolutely crazy. In other words, this form of relativism is an implication of a certain form of rule consequentialism; if one has good reasons to accept the latter, the former is also acceptable. Given the space limit, I will not fully evaluate the plausibility of this position. For the sake of argument, let us suppose the reviewer is right about the connection between a certain form of rule consequentialism and the relevant relativism. There are, however, at least two things that can be said here. First, many people are not rule consequentialists, so the argument mentioned above provides no reason for them to accept the relevant relativism. Second, for people who do find this form of relativism rather implausible, the reviewer's point actually provides a reason to reject the corresponding form of rule consequentialism, rather than a reason to accept the relevant relativism.

Before we move on, one additional remark should be made. In the thought experiment discussed above, I focused on the difference in terms of cognitive capacities. I argued that this kind of difference, combined with Kagan's hierarchicalism and practical realism, leads to implausible moral status relativism. One may point out that practical realism also tells us to attend to motivational limitations (see Kagan 2019: 285). The idea might be that if we are simply incapable of being motivated to conform to a set of rules that would direct us to assign different unimpaired people different moral statuses, then, according to practical realism, we should adopt a set of rules that would assign the same moral status to all unimpaired adults. However, this strategy does not avoid the problem of moral status relativism. For example, suppose that the aforementioned alien species is also capable of being motivated to conform to a set of rules that would direct them to assign different moral statuses to different unimpaired human adults. In this case, too, humans and the aliens will have different versions of the limited hierarchy, and thus the problem of moral status relativism persists.

4.2.2. *Moral status and time*

Besides variations of moral status relative to species, there might also be variations relative to time. Specifically, future human beings may have moral status systems that are different to ours for a reason analogous to the one that holds for the alien species described in the previous section. Indeed, owing to technological developments, it is plausible to suppose that future humans will have enhanced abilities to discern further nuances in an individual's relevant psychological capacities. Furthermore, they may be able to deal with more complicated moral status systems than ours. In short, we may become beings like the aliens discussed above. Just as the alien has a different literally true moral status system from the one adopted by current humans, future human beings will have a different moral status system as well. Again, we fall into moral status relativism.¹⁹

Moreover, even if we were to accept practical realism regardless of the moral status relativism it entails, the problem of normal variation would reappear. As our cognitive abilities become more advanced, at some point, we will become sensitive to the difference between different unimpaired adults' psychological capacities. Future humans will assign different moral statuses to different unimpaired adults, and the problem of normal variation will reappear. Consequently, even if we put moral relativism aside, practical realism, at best, postpones, rather than solves, the problem. In other words, Kagan's hierarchicalism has the defect of being *predictably problematic*.

People such as Kagan might suggest that it is possible that when people have such great cognitive capacities, they may find that normal variation is acceptable. After all, the rejection of this result is derived from the morality we have today. As time passes, and assuming certain changes in our psychological capacities, we may acquire different moral intuitions. This suggestion has some plausibility as a possibility, but there is no compelling evidence that this possibility will likely be realized. Given how strong our intuition is that normal variation is unacceptable, pointing out the mere possibility that such intuition might change does not weaken my argument. Moreover, this reply cannot solve the problem of moral status relativism.

In short, defenders of Kagan's hierarchicalism cannot embrace practical realism as a solution to the problem of normal variation. First, endorsing those two views leads to moral status relativism, which is implausible. Second, practical realism, as a solution, is

¹⁹Again, the consideration of motivational limitations does not help.

predictably problematic insofar as it is highly likely to become problematic in the future if humans gain a more fine-grained moral status system than ours, a system that will make the problem of normal variation reappear. Crucially, without practical realism (and practical concern), Kagan's hierarchicalism lacks a tenable solution to the problem of normal variation. This leaves the view at odds with the widely held belief that all unimpaired adults have the same moral status.

5. Conclusion

In this paper, I have provided a twofold critique of Kagan's hierarchicalism.²⁰ First, I addressed the primary argument that favors Kagan's view over unitarianism, namely, the argument from distribution. I suggested that the key underlying intuition – K-intuition – is, in all likelihood, rooted in the bias of the powerful. By challenging the validity of K-intuition, I have endeavored to undermine the force of the argument from distribution, thereby reducing the appeal of Kagan's stance.

In the second part of the paper, I argued that there is a good pro tanto reason to reject Kagan's hierarchicalism on the grounds that he fails to explain away the counter-intuitive result of his theory in the case of normal variation. Prima facie, hierarchicalism implies that even among cognitively unimpaired adults, there exists variation in moral status. This result seems counterintuitive and consequently casts doubt on hierarchicalism. I argued that Kagan's two attempts to solve this problem are unsuccessful. On the one hand, the first solution Kagan proposes, practical concern, is obviously unsatisfactory because it contains the unacceptable presupposition that there are indeed differences among ordinary unimpaired people's moral statuses. On the other hand, the alternative solution that Kagan puts forth, namely practical realism, is predictably problematic and leads to the unsatisfactory result of moral status relativism. Consequently, without a plausible solution, Kagan's hierarchicalism remains vulnerable to the significant challenge posed by the problem of normal variation.

Acknowledgements. I would like to thank Ann Victoria Levey, Peter Vallentyne, Ryan Kulesa, Joel Vall Thomas, Se Yong Bae, Shan Jiang, Alok Tiwari, David DeGrazia, and two anonymous reviewers for their helpful comments on earlier versions of this paper.

References

- Anderson, Elizabeth.** 2014. Social Movements, Experiments in Living, and Moral Progress: Case Studies from Britain's Abolition of Slavery. *The Lindley Lecture at the University of Kansas* (Lawrence, KS: University of Kansas, Department of Philosophy).
- Anderson, Elizabeth.** 2021. How to Be a Pragmatist, in *The Routledge Handbook of Practical Reason*, ed. by Ruth Chang and Kurt Sylvan (Abingdon: Routledge): 83–94.
- Beauchamp, Tom L.** (ed.) 1998. *An Enquiry Concerning the Principles of Morals* (Oxford: Oxford University Press).
- Brouwer, Huub, and Willem van der Deijl.** 2020. All Animals are Equal, but Some More than Others? *Journal of Moral Philosophy* 17 (3): 342–57.

²⁰My objection is new when compared with existing objections. Critics primarily target Kagan's view of modal capacity and the value of the wellbeing argument. On the one hand, DeGrazia (2015), McMahan (2015), and Lloyd (2021) question Kagan's idea that modal capacities can raise one's moral status. On the other hand, Lloyd (2021) provides a detailed and convincing argument against another argument, i.e., the value of the wellbeing argument.

- DeGrazia, David.** 2015. Modal Personhood and Moral Status: A Reply to Kagan's Proposal. *Journal of Applied Philosophy* 33 (1): 22–25.
- Dewey, John and Tufts, James.** 1932. *Ethics* ed. Jo Ann Boydston. Reprinted *The Later Works, 1925–1953*. (Carbondale: Southern Illinois University Press, 1981).
- Kagan, Shelly.** 2016. What's Wrong with Speciesism? *Journal of Applied Philosophy* 33 (1): 1–21.
- Kagan, Shelly.** 2019. *How to Count Animals, More or Less* (Oxford: Oxford University Press).
- Kazez, Jean.** 2010. *Animalkind: What We Owe to Animals* (Oxford: Wiley-Blackwell).
- Lloyd, Harry R.** 2021. Moral Status, Luck, and Modal Capacities: Debating Shelly Kagan. *Journal of Applied Philosophy* 38 (2): 273–87.
- McMahan, Jeff.** 1996. Cognitive Disability, Misfortune, and Justice. *Philosophy & Public Affairs* 25 (1): 3–35.
- McMahan, Jeff.** 2015. On 'Modal Personism'. *Journal of Applied Philosophy* 33 (1): 26–30.
- Sebo, Jeff.** 2021. How to Count Animals, More or Less, by Shelly Kagan. *Mind* 130 (518): 689–97.
- Singer, Peter.** 1990. *Animal Liberation*, 2nd edn (New York: New York Review of Books).
- Singer, Peter.** 2009. *Animal Liberation: The Definitive Classic of the Animal Movement*, Reissue edn (New York: Harper Perennial Modern Classics).
- Vallentyne, Peter.** 2005. Of Mice and Men: Equality and Animals. *The Journal of Ethics* 9 (3–4): 403–33.