



How to Measure Capacity for Welfare and Moral Status

Jason Schukraft

June 1, 2020

Executive Summary

An animal's capacity for welfare is how good or bad its life can go. An animal's moral status is the degree to which an animal's experiences or interests matter morally. It's plausible that animals differ in their capacity for welfare and/or their moral status. These differences could affect the way we ought to allocate resources across interventions and/or cause areas. Unfortunately, measuring capacity for welfare and moral status is tremendously difficult.

When donors or researchers choose to focus on cause areas or interventions that target certain species rather than others, they are often implicitly making judgments about the comparative value of different animals (including humans). Without a model for quantifying differences in comparative value, such judgments are apt to be guided by imperfect and likely unreliable heuristics.

There are two non-exclusive methods we might employ to measure capacity for welfare and moral status. The first method is to survey various experts about what sort of tradeoffs among animals they would endorse. This approach is relatively simple and cheap, but it relies on the assumption that intuitions about moral tradeoffs reliably track the moral truth. This assumption looks dubious. Intuitive judgments of this kind are often sensitive to non-evidential factors. Deep-rooted, widespread speciesism is likely to prejudice responses.

The second method is more time-consuming and complex but potentially more objective. The method proceeds in three steps. The first step is to canvass the relevant philosophical literature to generate a relatively theory-neutral list of characteristics that might contribute to capacity for welfare or moral status. The second step is to find empirically measurable proxies for those characteristics and weight the proxies by their relative importance. The third step is to canvass the relevant scientific literature to

score different animals of interest according to the features identified in the second step. Estimates of uncertainty would be made at each step, and a sensitivity analysis would help identify areas of high information value. I estimate that such a project would require between five thousand and seven thousand person-hours to complete.

Introduction and Context

This post is the second in [Rethink Priorities'](#) series about comparing capacity for welfare and moral status across species. The primary goal of this series is to improve the way resources are allocated within the effective animal advocacy movement. A secondary goal is to improve the allocation of resources between human-focused cause areas and nonhuman-animal-focused cause areas. In the [first post](#) I lay the conceptual framework for the rest of the series, outlining different theories of welfare and moral status and the relationship between the two. In this, the second entry in the series, I present and evaluate two methodologies for measuring and comparing capacity for welfare and moral status. In the [third post](#) I explain what the subjective experience of time is, why it matters, and why it's plausible that there are morally significant differences in the subjective experience of time across species. In the [fourth post](#) in the series, I explore critical flicker-fusion frequency as a potential proxy for the subjective experience of time.

The Measurement Problem

Humans exploit a huge variety of animals. On an annual basis, humans slaughter about [290 million frogs](#), [480 million goats](#), [2.9 billion snails](#), [3 billion ducks](#), [22 billion cochineal bugs](#), [69 billion chickens](#), [300 billion crustaceans](#), and nearly [a trillion commercially caught fish](#).¹ At any given time, humans confine about [251 million sheep](#), [265 million cows](#), [7.5 billion hens](#), and more than



[1.4 trillion bees](#) to produce wool, milk, eggs, and honey. Counting somewhat conservatively, humans exploit at least 33 orders of animals, across 13 classes and 6 phyla.¹ The effective animal advocacy (EAA) movement has limited resources, and it must choose how to allocate these scarce resources among these different animals, most of whom are treated miserably by humans. Since we can't (yet) help all these animals, we must decide which animals to prioritize.²

In the [first entry](#) in this series, I argued that there are good reasons to think that animals differ in their [capacity for welfare](#) and/or their [moral status](#). I also claimed that these differences could significantly affect the way we ought to allocate resources across cause areas and interventions. Of course, these differences can only affect our allocative decision-making if we know about them and know, at least roughly, their magnitudes. Hence, it is important that we devise a method for reliably measuring capacity for welfare and moral status and comparing them across species.³

Methods for measuring capacity for welfare and moral status can be assessed across a number of important criteria. The method must be *valid* and *accurate*—that is, it must actually track capacity for welfare and moral status and be sensitive to differences in capacity for welfare and moral status. Ideally, the method would be *applicable across species*—that is, it would be accurate and valid with respect to phylogenetically distant animals occupying radically different ecological niches.⁴ Ideally, the method would be *sensitive to moral uncertainty*—that is, rather than assume a particular normative framework, the method would allow one to input a variety of plausible axiological assumptions and observe how changing the assumptions changes the outputs of the final model. The *practical feasibility* of the method must also be considered. How simple is the method to execute and use when finished? How much would executing the method cost? How likely is it that an attempt to execute the method ends in failure?⁵

In this post I consider two methods for measuring and comparing capacity for welfare and moral status: (1) a holistic approach, in which relevant experts employ their normative and biological

expertise to make all-things-considered estimates of the appropriate tradeoffs between different lives, experiences, or interests, and (2) an atomistic approach, in which we identify empirical proxies for morally salient features, then let our best scientific understanding of the degree to which different animals possess those features guide our estimates of comparative moral value. The two approaches are not in principle mutually exclusive. One could in theory adopt both approaches, then let one's final estimates be conditioned by a weighted [reflective equilibrium](#) between the two.

I argue that the atomistic approach is the more difficult but ultimately the more accurate method. Thus, any reflective equilibrium between the two approaches ought to be weighted more heavily toward the atomistic rather than the holistic approach. Nonetheless, the atomistic approach faces serious complications along at least three dimensions: identifying empirically-measurable proxies for the morally salient features, comparing those proxies across phylogenetically distant animals, and incorporating differential performance on those features into a unified, common metric weighted by the importance of the features.

Of course, we shouldn't expect that we'll ever be able to pinpoint an animal's *precise* capacity for welfare or moral status. As I detail later in this series, there is a tremendous amount of empirical uncertainty about the extent to which different animals display different morally relevant traits and features. And even if the empirical uncertainty could be resolved, the philosophical uncertainty would likely remain.⁶ Thus, our best methodology executed as well as we can will still deliver only ranges of values, and it's difficult to say in advance how wide those ranges will be. Attempting to measure capacity for welfare and moral status will help us identify our degree of uncertainty regarding these issues. Merely knowing the extent of our uncertainty could plausibly improve our decision-making process.⁷

1 See [this spreadsheet](#) for details. By my count, every order in the spreadsheet is exploited in numbers greater than ~50 million individuals per year.

2 Humans also indirectly affect many wild animals, and many wild animals suffer independent of any human interference. In this series I focus primarily on animals that humans exploit directly, most of which are farmed. Because the goal of the project is to improve the way resources are allocated across interventions, it makes sense at this time to focus on animals that are directly exploited. As the effective animal advocacy movement identifies more interventions to aid wild animals, we will want to include those animals in our measures of comparative moral value.

3 See [Budolfson & Spears 2019](#) for more on the measurement problem.

4 A method that was more limited (to say farmed land vertebrates) could still be useful even if less than ideal.

5 This list is adapted from [Browning 2020](#). Her focus is on measuring *realized* welfare, but I think the desiderata apply equally well to measuring capacity for welfare and moral status.

6 I'm not claiming here that philosophical disputes are in principle irresolvable, just that they are usually much less tractable than empirical questions.

7 For instance, we may want to adopt a precautionary principle ([Birch 2017](#)) in the face of large uncertainty.

The Holistic Approach

According to what I'm calling the *holistic approach* to measuring capacity for welfare and moral status, the best way to estimate moral status and capacity for welfare is to think holistically about the comparative value of different sorts of animals. The approach is *holistic* because it starts at the question we are trying to answer rather than trying to decompose the question into constituent parts. The holistic approach elicits all-things-considered judgments about the relative value of different animals, and there is no underlying framework which determines which considerations ought to bear on the final judgments.

Insofar as there is a currently endorsed method for adjudicating disputes about the comparative value of different animals, the holistic approach seems to be the preferred method. As far as I can tell, in most organizations decisions about the comparative value of different animals are governed by intuitive judgments rather than tables and spreadsheets.⁸ This is probably for the best. Capacity for welfare and moral status are complicated topics; they don't lend themselves to easy formulae. Explicit numerical representations are apt to gloss over important complexity, and simple quantitative models are unlikely to outperform all-things-considered judgments by domain experts. As we'll see below, I estimate that executing the atomistic approach to the rough specifications I outline would require about six thousand person-hours, with little action-guiding payoff until at least midway through the project. Since most organizations don't have three to six thousand hours to think about these issues, the holistic approach is an acceptable short-term stopgap. In the medium- to long-term, the only way to ensure that we are efficiently allocating resources across different groups of animals is to invest the time and money necessary to thoroughly study moral status and capacity for welfare.

Tradeoffs and preferences

One way to attempt to measure comparative moral value is by directly judging what sort of tradeoffs between different species

would be appropriate. The tradeoffs might be couched in terms of lives: we might wonder how many salmon we ought to be willing to let die in order to save one thousand turkeys. The tradeoffs might be couched in terms of experiences: we might wonder how many minutes of suffering we ought to be willing to let a lobster endure in order to alleviate one hundred minutes of frog suffering. Or the tradeoffs might be couched in terms of interests: assuming pigs and chickens have an equally strong interest in avoiding extreme confinement, we might wonder how many hens we ought to be willing to forgo freeing in order to liberate ten sows.

Another approach is to couch the tradeoffs in terms of what species one would prefer to be. For example, Peter Singer allows that "it would not necessarily be speciesist to rank the value of different lives in some hierarchical ordering. *How* we should go about doing this is another question, and I have nothing better to offer than the imaginative reconstruction of what it would be like to be a different kind of being" (Singer 2011: 91).⁹ He suggests that "If it is true that we can make sense of the choice between existence as a horse and existence as a human, then – whichever way the choice would go – we can make sense of the idea that the life of one kind of animal possesses greater value than the life of another; and if this is so, then the claim that the life of every being has equal value is on very weak ground" (Singer 2011: 91). With a suitably large and diverse sample of matched pairs, we could create an ordered ranking.

We could also ask how many days of one's human life one would be willing to forgo to experience some duration of time as another species. This approach would allow us to assign cardinal numbers to the value of animal lives. Shelly Kagan imagines such an approach. He writes, "The average human life span is about 79 years, or more than 28,000 days. Divided by ten thousand that's still more than 2.8 days. If, like me, you wouldn't give up even a *single* day as a person for an entire extra lifetime as a fly, then you agree that the welfare to be had within a fly's life is less than one ten thousandth the welfare to be found in a person's life" (Kagan 2019: 90, fn 5).¹⁰ Again, by considering one's preferred tradeoffs across a large and diverse sample of different animals, we could

⁸ For a recent example to the contrary, see Founders Pledge's [report](#) comparing the value of donations to The Humane League to the value of donations to the Against Malaria Foundation. For another potential example to the contrary, see Charity Entrepreneurship's [weighted welfare index](#). Note that although the CE index is meant to improve the way resources are allocated across species, it does not explicitly address moral status or capacity for welfare.

⁹ He adds, "Some comparisons may be too difficult. We may have to say that we have not the slightest idea whether it would be better to be a fish or a snake; but then, we do not very often find ourselves forced to choose between killing a fish or a snake. Other comparisons might not be so difficult. In general, it does seem that the more highly developed the mental life of the being, the greater the degree of self-awareness and rationality and the broader the range of possible experiences, the more one would prefer that kind of life, if one were choosing between it and a being at a lower level of awareness" (Singer 2011: 92).

¹⁰ Note that Kagan's presentation is a bit misleading here. Comparing the welfare of a human's lifetime with the welfare of a fly's lifetime is a comparison in *diachronic* welfare, that is, welfare over time. But humans live much longer than flies; thus they have much longer to amass welfare. So it's not really a fair comparison. Fruit flies only live about 30 days. If one refused to forgo a single day of human life for an extra lifetime (30 days) as a fly, we should infer that a day in the life of a typical fly contains less than one thirtieth the welfare of a day in the life of a typical human.

begin to construct a hierarchy of comparative moral value.

Survey data

There is already a wealth of existing survey data about attitudes to animals, and some of this data can be repurposed to infer the general public's positions on the morally appropriate tradeoffs among species. The Animal Attitudes Scale (AAS) has been in use since 2002.¹¹ The AAS and its variants¹² ask respondents to agree or disagree (on a five point scale) with twenty-eight statements such as *It is morally wrong to hunt wild animals just for sport* and *Breeding animals for their skins is a legitimate use of animals*. Some of the questions refer to specific animals, such as *It is morally wrong to eat chicken and fish* and *A human has no right to use a horse as a means of transportation (riding) or entertainment (racing)*. Comparing such responses reveals rough differences in attitudes toward different animals, but it does not reveal the degree to which some animals are valued more than others.

A new scale, the Animal Purpose Questionnaire (APQ), has recently been devised to offer “a more differentiated measure of attitudes to animal use across a variety of settings” (Bradley et al 2020: 1). The APQ asks respondents the extent to which they agree (on a five point scale) that it's permissible for animals to be killed for different purposes. In total the APQ asks about sixteen animals¹³ and five uses,¹⁴ though the survey is designed so that respondents aren't asked about all animals and all uses. Generalizing across respondents and usage categories, Bradley et al. 2020 find that respondents tend to value monkeys, badgers, tree shrews, chimpanzees, dogs, dolphins and parrots more highly than rats, mice, pigs, octopuses, chickens, zebrafish, carp, and pigeons. Again, though, the scale cannot pinpoint the exact extent to which some animals are valued more than others.

Another recent survey, Miralles et al. 2019, asked respondents about their relative levels of empathy and compassion for animals of different species. Each respondent was asked to view pictures of two animals of different species. For the empathy questions, respondents chose the animal for which they felt like they were “better able to understand the feelings or emotions of.” For the compassion questions, respondents chose which animal they would save if both were in danger of death. Both empathy and compassion decreased with increasing phylogenetic distance

from humans. However, once again, this survey methodology does not allow us to infer the exact numerical tradeoffs between animals that respondents would endorse.

The general public has occasionally been surveyed about specific tradeoffs. For instance, in March 2019 Scott Alexander ran a [small survey](#) (n=50) asking respondents to estimate the relative value of nonhuman animal lives in comparison to a human life. The median respondent in his survey estimated that a single human is as valuable as 4,000 lobsters, 500 chickens, 50 cows, 35 pigs, 7 elephants, or 5 chimpanzees. Shortly thereafter, a commenter called Tibbar [attempted to replicate Alexander's survey](#) with a larger pool of respondents (n=263). The results were strikingly different, with Tibbar's respondents ranking the relative value of a human life much lower than Alexander's respondents. According to the median respondent in Tibbar's survey, a human life is as valuable as 60 lobsters, 25 chickens, 5 pigs, 3 cows, and 2 chimpanzees. (Elephants scored as highly as humans.)

Rethink Priorities offered to investigate the discrepancy between Alexander's and Tibbar's results. We launched a new, larger survey (n=490) and found enormous variance in the value assigned to different animals. Many respondents assigned each animal a value equal to humans, and many respondents did essentially the opposite—indicating that human life was incommensurable with or infinitely more valuable than nonhuman animal life. In between these positions there was an extreme range, with some respondents assigning a value to each animal *nearly* equal with humans and other respondents assigning nonhuman animals quadrillions times lower moral value than a single human. Such variegated data posed many interpretative challenges, but ultimately we concluded there were two natural ways to analyze the data, one of which supported Alexander's high values and one of which supported Tibbar's lower figures. Alexander [reported on our findings](#) in May 2019. A full write-up from Rethink Priorities is forthcoming.

Such surveys are not limited to popular blogs. For example, in a 2007 phone survey of one thousand Americans, Oklahoma State agricultural economists Bailey Norwood and Jason Lusk asked respondents to agree or disagree with the following statement: “If a new technology were created that could either eliminate the suffering of 1 human or the suffering of X farm animals, it should be used to eliminate the suffering of the 1 human.” The variable X was randomly set to 1, 10, 50, 100, 500, 1,000, 5,000, or 10,000.

11 See [Wuensch, Jenkins, & Poteat 2002](#).

12 See [Herzog, Grayson, & McCord 2015](#) for a shorter version of the scale.

13 Mice, rats, rabbits, pigs, monkeys, octopuses, chickens, badgers, zebrafish, tree shrews, dogs, dolphins, parrots, chimpanzees, badgers, and pigeons.

14 The specific uses are medical research, basic science research, food production, pest control, and “other.” Note that specifying specific uses probably introduces many confounding influences to the responses.

Extrapolating from the results, Norwood and Lusk concluded that the average American believes that the suffering of one human is equivalent to the suffering of about 11,500 farm animals.¹⁵

A more recent survey of this type is reported in [Weathers et al. 2020](#). Respondents were asked to compare the suffering of cows, pigs, and chickens via a series of tradeoff questions. For example, a respondent might be asked to compare two hypothetical programs, the first of which would prevent one thousand cows from contracting an illness that causes rapid death and the second of which would prevent X chickens from contracting a similar illness that causes rapid death. Respondents were then required to select the lowest value of X for which the second program would produce the greater overall reduction in suffering, with possible values for X ranging from one to one million.¹⁶ According to the authors, “Approximately 39.9% of participants valued cattle more than chickens, and a similar proportion (38.8%) valued pigs more than chickens” ([Weathers et al. 2020](#): 4).¹⁷

Few of these survey designs are ideal and none is perfect. Nevertheless, the data presented here do yield at least one tentative conclusion: many people are comfortable endorsing a hierarchy of moral value. It appears to be a commonly—if not universally—accepted view that some animals are more valuable than others, even if there is disagreement as to the extent of the differences. What remains to be seen is whether or not this position is justified. In the following section, I present some evidence that lay intuitions about comparative moral value should not be trusted.

The problem with appeals to intuition

One initial concern about appeals to intuition in this domain is the general zoological ignorance of the intuiting public. As I noted above, humans directly exploit at least 33 orders of animals across 13 classes and 6 phyla. The average person simply doesn’t know much detail about the lives of, say, goats, geese, carp, catfish, earthworms, silkworms, snails, or squid.¹⁸ But without detailed knowledge of the characteristics of the species under comparison, it’s hard to see what could justify judgments of comparative moral worth. However, I want to set this worry aside. I assume that if we adopted the holistic approach, we would only care to elicit the intuitions of qualified experts.¹⁹ My worry is that even the intuitions of zoological experts will be unreliable.

The holistic approach is driven by all-things-considered judgments about which species one ought to prefer to be and which tradeoffs between species are morally appropriate. By their very nature, the origin of these judgments is somewhat opaque. The judgments are not the product of a clearly delineated algorithm or decision tree. They don’t imply whether variation in moral value is due to differences in moral status or differences in capacity for welfare (or both). They certainly don’t say which particular features are driving the difference in moral value. In many ways this is a feature, not a bug: developing a formula for calculating moral value is difficult, and without rigorous, protracted investigation, such a formula is unlikely to outperform rapid intuitive judgment. But the speed of intuitive judgment comes at a price: when the origin of a judgment is opaque, it’s easier for unwanted influences to creep in without one’s knowledge.²⁰

¹⁵ The survey was conducted for the [American Farm Bureau Federation](#) and is unfortunately no longer available online. Norwood and Lusk discuss the survey in their 2011 book [Compassion, by the Pound: The Economics of Farm Animal Welfare](#) (pp. 171-172). A popularization of the survey appeared in *Reason* magazine under the title “[You=11,500 Sheep](#).”

¹⁶ The survey utilized a multiple-choice format, so respondents were not able to input any number they wanted for X. For all questions, the first program affected one thousand animals and the possible values for X for the second program were 1, 500, 1001, 2000, 5000, 10000, 100000, and 1000000.

¹⁷ Also of note: “Nearly one-third (30%) of respondents reported that they believed animal suffering should be taken into account to a degree equal to or above human suffering” (1).

¹⁸ As just one example of the extent of this ignorance, I hypothesize that few members of the general public would guess that snails are more closely related to squid than earthworms are to silkworms.

¹⁹ It may still be useful to survey the lay public. Such surveys may help us identify biases that influence ours and others’ judgments. It’s also possible that a wide enough survey may reveal a latent ‘wisdom of the crowd,’ which would allow us to extract a useful signal from the random noise of our unreliable, arbitrary intuitions.

²⁰ See [Cuddington et al. 2013](#) for more on the tradeoff between the speed and opacity of expert judgment: “While the development of a rule may take some time, expert opinion can be accessed rapidly in most cases, and in some cases is the only information available (O’Neill et al. 2008). However, the role of theory and the assumptions behind expert opinion and rules of thumb are rarely transparent, so there may be little potential for evaluating the assumptions that support models of this sort. Expert opinions inevitably are divergent (e.g., Czembor et al. 2011), although there may be techniques for building consensus among a group of experts (e.g., Delphi technique, Rowe and Wright 1999). It is also possible that the rules of thumb or expert opinion do not include adequate concepts of scale and uncertainty (e.g., Burgman 2005) that are a requirement for appropriate management under global change. However, when the main requirement is that a decision be made extremely quickly with very limited data, expert opinion or rule-based models have a clear time advantage over other types of models” (3).

There is already a large literature which demonstrates that uncalibrated intuitions are often sensitive to non-evidential factors.²¹ Since intuitions about the comparative moral value of nonhuman animals are not amenable to independent calibration, these intuitions are almost certainly influenced to some degree by factors that are morally irrelevant. So I think there is good reason *in general* to worry that unwanted considerations unduly sway one's intuitions about the value of nonhuman animals. To compound this general worry, there are reasons to think that, in the specific case at hand, irrelevant factors are likely to unconsciously taint our reasoning.

There is ample evidence that humans tend to value large mammals, especially those with big eyes or human-like characteristics, over other animals. Animals with fur are preferred to those with scales; animals with two or four limbs are preferred to those with six, eight, or ten limbs. Animals deemed to be aggressive, dirty, or dangerous are perceived negatively. Companion animals and pets attract more sympathy than comparable farmed or wild animals.²² These factors, and many others, will plausibly influence our reactions to thought experiments.

Consider, for instance, Singer's invitation to assess whether we would prefer to experience life as one species rather than another. The goal of the exercise is to use our imaginative faculties to estimate the capacity for welfare that different animals possess.²³ In addition to the above biases, such thought experiments may be swayed by personal considerations that have nothing to do with capacity for welfare. Perhaps I would prefer to be a marlin rather than a chimpanzee because I like to swim. Perhaps I would rather be a gecko than a polar bear because I dislike cold climes. Perhaps I'd prefer to be a penguin rather than a snake because I know penguins engender more sympathy than snakes.

Similar concerns haunt Kagan's invitation to consider how much of one's human life one would sacrifice to gain an extra lifetime as a member of another species. Perhaps I would gladly sacrifice a year of my human life for an extra lifetime as a sparrow because the novelty of unaided flight intrigues me. Of course, in carefully presenting the thought experiment we would stipulate that such personal preferences are irrelevant and ought to be bracketed when considering the tradeoff. But it's unclear exactly which per-

sonal preferences *are* irrelevant, and even if we were confident in our delineation of relevant versus irrelevant preferences, it's an open question how successfully we can bracket the irrelevant influences by stipulation.²⁴ Moreover, it's also a bit unclear how Kagan's thought experiments could tell us much about moral status (as opposed to capacity for welfare). Moral status presumably makes no intrinsic contribution to phenomenology, so if two animals had the same capacity for welfare but different moral statuses, it's unclear why I should be willing to give up more time for an extra life as the animal with the higher status.

The surveys in the previous section sometimes report results that are best understood if we accept that irrelevant factors often drive our intuitive responses. For instance, in the APQ, respondents typically don't object to killing mice and rats for medical research, basic science research, or pest control, but respondents *do* object to them being killed for food production. This is almost certainly because respondents find the idea of eating mice and rats unappealing, which, of course, is totally beside the point ([Bradley et al. 2020](#): 17-18).

Then there is the specter of speciesism, which in simple terms is a prejudice in favor of one's own species. Jeff Sebo warns that "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, will anchor us to an unacceptably conservative and speciesist moral theory" ([Sebo 2020](#): 6). Sebo adds that "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, that is, facts about the animal, rather than to external causes, that is, 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" ([Sebo 2020](#): 6). Sebo focuses his criticism on comparisons between humans and nonhumans. But what about comparisons restricted only to nonhuman animals?

The belief that humans are more valuable than nonhumans is

21 For an overview of the biases and heuristics literature, see [Kahneman 2011](#). For philosophical examples, see, among others, [Machery et al. 2004](#); [Swain, Alexander, & Weinberg 2008](#); [Buckwalter & Stich 2014](#); and [Costa et al. 2014](#). See [Schwitzgebel & Cushman 2012](#) for a series of experiments in which the moral judgments of professional philosophers were as sensitive to order effects as the judgments of non-philosophers. See [De Cruz 2015](#) §6 for a general discussion about whether these findings undermine the view that professional philosophers are 'expert intuiters.'

22 See [Serpell 2004](#), [Wynne 2007](#), and [Herzog 2010](#) for overviews.

23 Alternatively, the exercise might tell us something about average realized welfare rather than capacity for welfare. A species might have a huge capacity for welfare, but if in fact the members of that species tend to lead net-negative lives, we would never want to be a member of that species.

24 Similar concerns apply to measuring the relative value of human health outcomes, which is crucial for calculating QALYs.

often a manifestation of speciesism. But the idea that some non-human animals are worth more than some *other* nonhuman animals isn't *obviously* speciesist. If I judge that chickens are more valuable than trout, it's not obvious how such a judgment reflects a prejudice in favor of my own species. Nonetheless, the judgment might still be speciesist if humans are unjustifiably taken to be the standard against which nonhuman animals are compared, even when they are compared against each other. Similarity to humans might be the metric by which most people evaluate the comparative moral value of nonhuman animals, and if that is a speciesist criterion, then the comparisons will be tainted by speciesism.

The Atomistic Approach

All told, I think the above worries suggest that we should search for a more objective approach to measuring comparative moral value. Of course, no approach to measuring comparative moral value will be completely devoid of appeals to intuition and immune to the influence of speciesism. However, by standardizing the inputs to the model and tying those inputs to empirical data, we can make our intuitions explicit and public, the better to judge them. Such a model would hopefully reduce the extent to which we are swayed by non-evidential factors. It would also enable us to pinpoint the differences driving disagreements about judgments of comparative moral value, making such disagreements more productive.

The holistic approach to measuring capacity for welfare and moral status relies on all-things-considered subjective judgments. If such judgments are likely prone to errors that bias the process, we should be wary of the holistic approach. One way to help correct for these biases is to ground our judgments wherever possible in the hard facts of animal physiology, psychology, and ethology.²⁵ Despite Kagan's invitation to entertain various thought experiments, he concedes that "any particular judgments we might make about how one type of animal ranks in comparison to others will be subject to revision in light of further advances in empirical science. We may well discover that we have overestimated or underestimated the psychological capacities of any given type of animal. Unsurprisingly, then, any such ranking will remain tentative (and perhaps a bit rough as well). **But in principle, at least, a suitably informed ranking could be produced, and that ranking could then be improved upon as science reveals more about the details of animal psychology**" (Kagan 2019:

113-114, emphasis added).

The holistic approach begins with questions about the morally appropriate tradeoffs among species. An alternative is to first develop a rough system for adjudicating comparisons of capacity for welfare and moral status. What I'm calling the *atomistic approach* begins with the question 'What features and characteristics determine moral status and capacity for welfare?' then uses the answers to that question to determine the morally appropriate tradeoffs among species. The approach is *atomistic* because it decomposes the question of comparative moral value into discrete constituents (atoms) that are answered independently and then aggregated. This approach is necessarily more complicated, but the potential gains in accuracy may be worth it. In the following section, I outline what one such strategy for producing a suitably informed ranking might look like.

A rough guide to estimating moral status and capacity for welfare atomistically

Such an approach might proceed in three stages. The first stage would lay the conceptual framework for the project.²⁶ During this stage, one would specify which features are likely to determine capacity for welfare and moral status.²⁷ This stage would not require one to take a definitive stance on different theories of welfare and moral status, but it would require one to determine the implications of various plausible views. Because philosophical questions are notoriously difficult to resolve, we should be fairly uncertain about which theories of welfare and moral status are correct. Given this deep uncertainty, we should probably value interventions that are robust across a number of different plausible views. The goal of this stage would be twofold: (1) to generate a relatively theory-neutral list of characteristics that might contribute to capacity for welfare or moral status and (2) to understand the relative importance of the characteristics, weighted both by the importance of the characteristic within a given theory and by the probability that the theory is true.

The second stage would lay the methodological framework for the project. During this stage, one would operationalize the features and characteristics enumerated during the first stage into measurable proxies. This stage would require engagement with the empirical literature so as to know what in practice can be measured. But this stage would also require substantive theoret-

²⁵ Of course, which facts are available in the scientific literature might be influenced by bias. There is probably comparatively more information on traits that humans find interesting.

²⁶ In many respects, the [first post](#) in this series already plays this role. However, that post does not explicitly weight the features relevant for moral status or capacity for welfare.

²⁷ This stage might benefit from a survey of relevant philosophical experts.

ical reasoning, to judge which metrics are good proxies for the features we ultimately care about. A key goal of this stage would be to find measurable metrics that can be meaningfully compared across phylogenetically distant animals. The metrics must also be comparable in some sense to each other, so that they can be weighted against each other.

The third stage would be the simplest but the most time-intensive. First, we would select the animals to be investigated for the project (including the taxonomic rank at which to compare them). Because the goal of the project is to improve the way resources are allocated across interventions, it makes sense to select animals that humans directly exploit in very large numbers. Next, the relevant scientific literature would be systematically reviewed and organized, and the results compiled in a large database. The end-product might be a table consisting of ~30 features measured across ~30 orders of animals. [This template](#) provides an example framework.²⁸ (Note that the taxa and features are purely illustrative; they don't represent final judgments about what it would be worthwhile to investigate.) The database could either be used informally to guide and justify all-things-considered judgments, or it could be formalized into an algorithm that takes the table inputs and converts them into numerical estimates of comparative moral value across different animals. For the latter use, we ought also to conduct a sensitivity analysis and estimate our uncertainty for all the input parameters, so that we can identify where the value of new information is highest.

Such a project would be comparable in scope and structure to Rethink Priorities' [2019 work on invertebrate sentience](#). Based on that analogy, I estimate that measuring capacity for welfare and moral status in this way would require somewhere between five thousand and seven thousand person-hours. In the rest of this section, I discuss some theoretical and practical obstacles that would need to be overcome in order to adequately measure capacity for welfare and moral status via the atomistic approach I have outlined. The list is certainly non-exhaustive, but it should give a representative flavor of the difficulties the atomistic approach faces. The obstacles are presented in increasing order of seriousness.

Choosing taxonomic rank

It is difficult to choose the right level of generality at which to try to measure capacity for welfare and moral status. There are competing considerations at play in this decision. On the one hand, there is pressure to drill down to a fairly narrow taxon (genus or species, say). The higher up the taxonomic hierarchy one goes, the more phylogenetically diverse a taxon becomes. If a taxon becomes too diverse, then the fact that a particular animal within the taxon possesses some relevant feature doesn't guarantee that other animals within the taxon also possess the feature.

Recall from the [first post](#) in the series that, strictly speaking, capacity for welfare and moral status are properties of *individuals*.

		Mammalia										Aves			Reptilia	Amphibia	Acti			
		Artiodactyla			Lagomorpha	Perissodactyla		Carnivora			Rodentia	Galliformes	Anseriformes	Testudines	Anura	Clupeiformes	Cypriniformes	Cichliformes	Gadiformes	
		Bovidae	Carnivoridae	Suidae	Leporidae	Equidae	Canidae	Felidae	Mustelidae	Muridae	Phasianidae	Anatidae	Tringidae	Ranidae	Clupeidae	Euglenidae	Cyprinidae	Cichlidae	Gadidae	
Physiological	Neuron Count																			
	Presence of Nociceptors																			
	Connection of Nociceptors to Central Nervous System																			
	Presence of Endogenous Opioids																			
Cognitive	Episodic Memory																			
	Uncertainty Monitoring																			
	Abstract Thought																			
	Theory of Mind																			
Social	Tool Use																			
	Contextual Learning																			
	Average Size of Group																			
	Effects of Isolation																			
Affective	Roles an Individual Can Occupy																			
	Number of Referential Signals																			
	Social Learning																			
	Food Sharing																			
Experiential	Number of Emotional States																			
	Fear/Anxiety Behavior																			
	Cognitive Bias																			
	Emotional Contagion																			
Experiential	Social Buffering																			
	Response Slowing																			
Experiential	Critical Flicker Fusion																			
	Phenomenal Unity																			

²⁸ The cells are empty not only because the relevant scientific literature has not been surveyed but also because it is as yet unclear what *sort* of response the cells merit. Ideally, we want the cells to be comparable, even when the cells are reporting very different metrics, so it might make most sense to score each cell on some arbitrary scale (e.g., from 1 to 10). But standardizing the process looks really difficult. See the objections section for more discussion.

Obviously, it is not possible to investigate every individual animal that might be subject to some intervention in order to determine the individual animal's capacity for welfare or moral status. As Kagan puts it: "After all, it would hardly be feasible to expect us to undertake a detailed investigation of a given animal's specific psychological capacities each time we were going to interact with one. This makes it almost inevitable that in normal circumstances we will assign a given animal on the basis of its species (or, more likely still, on the basis of even larger, more general biological categories)" (Kagan 2019: 294). Moreover, we should expect that typical variation in capacity for welfare and moral status among members of the same species to be minimal. Generalizing to the level of species thus appears unproblematic.

Although generalizing to the level of species poses little theoretical difficulty, measuring capacity for welfare and moral status at that level is probably practically infeasible. There are hundreds of species that humans exploit in large numbers. Not only would it be difficult to investigate such a large number of animals, but the lower one goes in the taxonomic hierarchy, the less research is available that pertains to a given taxon. For all but the most commonly studied model organisms, it would be impossible to fill in the database at the level of species.

Moving a couple rungs up the taxonomic ladder to the rank of family improves the situation—but only slightly. There are approximately 50-60 families of animals that are directly exploited in large numbers. Moving up another rung in the ladder²⁹—to the rank of order—reduces the number of taxa to be investigated to a more manageable 33.³⁰ (Some of those animals, such as jellyfish, bivalves, and nematodes, might be thought to lack moral standing and thus could possibly be safely ignored, further reducing the final number.) The difference between investigating moral status and capacity for welfare at the rank of order rather than family could be as high as a thousand person-hours.

Unfortunately, measuring capacity for welfare and moral status at the rank of order may gloss over important differences among animals. To give just one example: humans and lemurs are both in the order *primates*. However, many people find it implausible that humans and lemurs have the same moral status or capacity for welfare.³¹ Thus, order may not be a fine-grained enough taxonomic rank to capture the relevant moral status facts.

29 Ignoring superfamilies, infraorders, and suborders, which not all orders have

30 See [this spreadsheet](#) for an overview.

31 One might think humans are unique and thus that this is a special case that says little about the general point. So here is a non-human example: among species in the *Carnivora* order, neuron count differs by a factor of 24 and brain mass differs by a factor of at least 58 (Jardim-Messeder et al. 2017).

32 Kagan believes there are only around six tiers of moral status (2019: 293).

33 See the [first post](#) in the series for details.

34 In practice, measuring neuron counts is actually anything but straightforward. See the work of [Suzanaerculano-Houzel](#) for many discussions of the var-

Ultimately, the choice of taxonomic rank in the project must be guided by a balance of considerations. Move too low, and the project balloons in size and the probability of finding relevant scientific studies for each taxon plummets. Move too high, and important, action-relevant information will be missed. Personally, I think order is probably the right rank at which to investigate the subject. In part this belief is driven by the view that if moral status admits of degrees, then moral status is discrete and organized into a relatively small number of levels.³² The smaller the number of tiers, the more sense it makes to investigate moral status at a higher taxonomic rank. It's less plausible that capacity for welfare is discrete, so if one thinks capacity for welfare is much more important than moral status in determining characteristic moral worth, then one probably ought to favor a more fine-grained investigation.

Finding measurable proxies

The atomistic approach recommends first canvassing the philosophical literature to ascertain which general features determine capacity for welfare and moral status. Across plausible philosophical views, there appears to be a rough consensus as to what sorts of general features are relevant for moral status and capacity for welfare. Examples include: intensity of valenced experiences, self-awareness, general intelligence, autonomy, long-term planning, communicative ability, affective complexity, self-governance, abstract thought, creativity, sociability, and normative evaluation.³³ However, it's one thing to identify general features that govern capacity for welfare and moral status. It's another matter entirely to find empirically measurable proxies for those features. If these features cannot be operationalized in a way that allows them to be measured, then the atomistic approach cannot succeed.

The gravity of this concern depends of course on which features one believes determine moral status and capacity for welfare. Some features are more amenable to operationalization than others. If one believed, for instance, that neuron count wholly determines moral status and capacity for welfare, then one would have a relatively straightforward method for measuring moral worth.³⁴ Sadly, the view that neurons determine moral status or

capacity for welfare appears rather unpromising. Neurons alone do not automatically generate conscious experience, and neurons are not themselves intrinsically morally valuable.³⁵ Larger animals need more neurons just to coordinate movement and autonomic functions. Larger animals also require more neurons to innervate their larger muscles,³⁶ and larger animals tend to process larger sensory fields to interact with their larger world, which requires a greater number of neurons just to process the data at the same level of complexity as a smaller animal would. Neuron counts alone do not tell us how the neurons are organized, how the neurons are used, or how many synaptic connections each neuron possesses. If neuron counts are worth investigating and comparing at all, it's only because they are themselves rough proxies for characteristics we care about. Perhaps neuron count correlates roughly with affective sophistication or intensity of valenced experience or general intelligence.

Unfortunately, many of these potentially morally important characteristics seem extremely difficult to operationalize, despite repeated attempts to do so. For instance, a feature as amorphous as *general intelligence* is unlikely to be captured by any single metric. The biologists Lesley Rogers and Gisela Kaplan put the point this way: "Intelligence is not an entity that can be measured by performance on just one task, nor can it be inferred from brain size, as we discuss below. Here it is worth noting that pigeons, tested on a task based on one problem taken from a standard IQ test for humans, which required them to recognize symbols rotated at different angles, surpassed humans in performance of the same task (Delius, 1987). Would we therefore rank them above us in intelligence? Obviously, the single criterion of assessment is an inadequate measure for intelligence in a broad sense. Although IQ tests have some degree of limited validity in terms of predictability of academic success in a given culture and class in humans (Sternberg, Grigorenko, and Bundy, 2001), there is in fact no scientifically acceptable way of measuring intelligence as a broad set of characteristics in humans, let alone in animals. Add to this the ambition of making comparisons of intelligence *across* species and it is easy to see how flawed such attempts would have to be" (Rogers & Kaplan 2004: 177-178).

Although I am not quite so pessimistic as Rogers and Kaplan, it's certainly worth acknowledging this difficulty at the outset of any attempt to measure general intelligence. Since many of the morally important features will be more akin to general intelligence than neuron count, we should expect the atomistic approach to

allocate many hundreds of person-hours to the task of identifying measurable proxies for morally salient features. This task will almost certainly require the collaboration of experts across multiple domains.

Comparing features across animals

Even if measurable proxies for the features we care about could be found, we would still have to compare those proxies across animals. Since humans exploit such a large and phylogenetically diverse range of animals, comparing the features is not going to be easy. Experiments exploring rodent intelligence likely look very different from experiments exploring eel intelligence. Experiments exploring self-control in chickens likely look very different from experiments exploring self-control in octopuses. Experiments exploring cow emotions likely look very different from experiments exploring fruit fly emotions. Experiments exploring the sociability of sheep likely look very different from experiments exploring sociability in honey bees. And so on. To properly compare results across experiments on different types of animals, some sort of normalization across studies is required.

Fortunately, there already exists a scientific discipline that aims to compare difficult-to-measure features across different animals. *Comparative cognition* is an interdisciplinary field at the intersection of animal psychology, neurology, ethology, and evolutionary biology. Any attempt to measure capacity for welfare or moral status across animals will almost certainly rely heavily on comparative cognition studies. There are prominent comparative cognition labs across the globe. Examples include [Cambridge's Comparative Cognition Lab](#), [University of Exeter's Centre for Research in Animal Behaviour](#), [Lund University's Cognitive Zoology Group](#), [Tufts University's Comparative Cognition Lab](#), [University of Helsinki's Comparative Mind Group](#), [Rochester Institute of Technology's Comparative Cognition & Perception Lab](#), and [IG-DORE's Interdisciplinary Research Group in Animal Behavioural Science](#).

In the last 15 years there has been a surge of interest in comparing species across a number of different interesting metrics. For example, [MacLean et al. 2014](#) compare "the cognitive performance of 567 individuals representing 36 species on two problem-solving tasks measuring self-control" (E2140). In addition to direct comparisons, there are also a number of meta-analyses that compile data from multiple studies to arrive at comparative conclusions.

ious complications.

35 We can imagine a lump of billions of neurons swirling around a laboratory jar with no capacity for welfare or moral standing. Conversely, we can imagine an alien or a computer program with zero neurons which nonetheless has a high moral status and capacity for welfare.

36 This does not necessarily give them greater precision in movement; insects on average have similar numbers of distinct muscles in total.

For example, [Cauchoix et al. 2018](#) “gathered 44 studies on individual performance of 25 species across six animal classes” in an effort to understand the evolution of cognition. Meanwhile, there has been a concomitant surge in theoretical discussions about how to compare features across species. For example [Weiss et al. 2019](#) outline a quantitative measure of social complexity that works across species and [Anderson & Andolphins 2014](#) develop a framework for studying emotions across species. These studies and others like them paint a promising picture of the potential of comparative cognition.

However, comparative cognition is a discipline still in its infancy. Because there may be a general bias toward non-null experimental results in the sciences, especially for relatively small and immature fields, we should be cautious about the conclusions of any one study ([Ioannidis 2005](#)). There is reason to think that [academic journals favor papers with surprising results](#) over papers which merely confirm the expected. Thus, there may be a publication bias in favor of animals doing surprising things. In the present case that might mean that comparative cognition studies which purport to demonstrate sophisticated cognitive abilities in nonhumans are overrepresented in the literature or that claims of comparability are exaggerated to gloss over undermining complications. Replication studies are, in general, under-rewarded in academia, so correcting for this overrepresentation and exaggeration may take years or even decades.

A recent (as yet unpublished) criticism of the field suggests comparative cognition research is biased because “(1) Phenomenon-based comparative cognition uses confirmatory research methods that are directionally biased, (2) In combination with a publication bias and a likely high rate of false discoveries, this bias suggests our literature contains many false positive findings, (3) This directional bias persists even with strong methodological criticism, and when researchers explicitly consider alternative explanations for the phenomena studied, (4) No formal method exists for generating and assessing theory-disconfirming evidence that could counter the biased positive evidence, (5) Ambiguity in definitions allow us as researchers to flexibly adjust our substantive claims depending on whether we are refuting criticism or selling the results, (6) The small size of comparative cognition as a research field perpetuates and reinforces points (1) - (5)” ([Farrar & Ostojic 2019](#): 4). Together, these points fa-

vor a healthy skepticism when new research analogizes behavior across species or ascribes surprisingly sophisticated cognitive abilities to nonhuman animals.³⁷

Some critics of comparative cognition worry that some questions the field pursues are premised on false assumptions. Daniel McShea complains of the “already fraught exercise of making comparisons across species lines,” wondering “how are we to compare the capabilities of, say, dusky titi monkeys with those of baboons? Dusky titis are smart about getting what they want, say, about the nuances of maintaining a pair bond. Baboons are also quite smart, but about different things, like navigating dominance hierarchies. Since the two species want such different things, since they are motivated to apply their non-affective capacities for such different purposes, one wonders whether it is even meaningful to ask which is smarter” ([McShea 2017](#): 7). Yasushi Kiyokawa and Michael Hennessy worry that “the variety of approaches [...] precludes any strict standardization of procedures. These factors and even the backgrounds of the researchers themselves will continue to promote ambiguity and differences of opinion” ([Kiyokawa & Hennessy 2018](#)). Writing of comparisons of pain states in different animals, Edgar Walters and Amanda Williams have noted “the difficulty in defining pain in a way that allows pain [...] to be recognized and compared across species, a task that is especially challenging for attempted comparisons of the conscious component of pain,” observing that “there is considerable uncertainty about which behavioural features, neural circuits, cell types and molecules to compare across taxa” ([Walters & Williams 2019](#): 6). In light of these difficulties, Lesley Rogers and Gisela Kaplan argue that “given our present state of knowledge of the needs and capabilities of classes of animals, let alone individual species, we feel, as biologists, that we first and foremost ought to guard against, or at least be very cautious about, the temptation of creating a scale of lesser or greater value of one species over another” ([Rogers & Kaplan 2004](#): 196).

Even the best case looks unpromising. Suppose we adopted the view that neuron count is a good proxy for moral status and capacity for welfare. At first blush, that seems like a relatively easy feature to measure and compare. But, as it happens, neurons aren’t all created equal because not all areas of the brain are equally important. Brain regions that, say, merely innervate muscles are plausibly less important to moral status and capacity for

37 On the other hand, there appears to be a longstanding and preexisting speciesist prejudice against attributing complex mental states to nonhuman animals. Given such a prejudice, scientists and the lay public may have been systematically underestimating the cognitive abilities of nonhuman animals for a long time. Today’s “surprising” results might just be the product of science finally beginning to overcome deep-rooted prejudice. And insofar as the prejudice persists, the competing forces of positive publication bias and speciesist prejudice might even approximately cancel each other out, leaving us with a literature that is largely reliable (though this is an unlikely outcome). In any event, the results captured in the comparative moral value database ought to be checked and updated on a continual basis. If there is a particular study that carries outsized weight in the final analysis, it may be worthwhile to fund a lab to replicate the study.

welfare than brain regions that, say, govern emotional responses.³⁸ Thus, across species, the number of neurons in certain brain regions may be more informative than overall neuron count. Two animals with the same overall number of neurons might differ in morally salient ways if those neurons are distributed differently across brain regions. But even when comparing the same comparably-sized brain regions across species, various cytoarchitectural differences, such as the extent of [cortical folding](#), [interneuronal distance](#), [axonal conduction velocity](#), [degree of myelination](#), and [synaptic transmission speed](#), could plausibly be more important still. To properly compare neurons, we need to know where they are located and how they are connected to each other. In many ways, the foregoing is an argument against taking neuron count to be a good proxy of moral status or capacity for welfare. But I hope it also serves to illustrate the general challenge of comparing even relatively simple physiological features, to say nothing of more complex, amorphous features such as general intelligence or affective sophistication.

All told, these worries suggest that comparing morally relevant features across phylogenetically distant animals will be fraught with theoretical and practical challenges. It will not always be clear which results are truly comparable, and as such, many difficult judgment calls will be required. It's possible that these subjective judgment calls will be so numerous and so inescapable that any pretense of objectivity will be lost. In that case, the best method for improving our ability to measure capacity for welfare and moral status might be to fund more rigorous work at various comparative cognition labs so that better procedural standards for comparison can be developed.

Weighting the features

After canvassing the philosophical literature to find characteristics that contribute to moral status and capacity for welfare and canvassing the scientific literature to find measurable proxies for those characteristics, we will have a general list of features to investigate.³⁹ However, there is good cause to anticipate that not all of those features will be equally important. There are at least three reasons to think that the features will need to be weighted.

The first reason is that within a given theory of moral status or capacity for welfare, some features are more significant than oth-

38 See [Dicke & Roth 2016](#) for more on the importance of cortical neurons, neuron packing density, interneuronal distance and axonal conduction velocity.

39 I tentatively estimate we'll end up with about 30 features on the list.

40 Examples of agential features include self-awareness, self-control, number of behavior types, executive functions, long-term planning, and capacity for moral responsibility.

41 Examples of physiological features include neuron count, presence of nociceptors, connection of nociceptors to central nervous system, and presence of endogenous opioids.

ers. This difference could be due to value pluralism. Objective list theories of welfare include among the list of intrinsic goods items such as happiness, virtue, wisdom, friendship, knowledge, and love. These items need not contribute to welfare equally. Even for a value monist view like hedonism, different features ought to be assigned different weights. There is no single proxy that perfectly captures capacity for pain and pleasure. A hedonist might think that self-awareness, linguistic sophistication, affective complexity, sociability, and long-term planning all influence the range and intensity of possible pleasures and pains a creature can experience. But it would be quite surprising if such diverse features contributed equally.

The second reason is that some features will be relevant to more theories than others, and in a relatively theory-neutral framework, the features ought to be weighted according to how many theories they are relevant to. Since the goal is to develop interventions that are robust in the face of our moral uncertainty, we should probably pay more attention to features that are salient across a spectrum of theories. For instance, although hedonism holds that pains and pleasures are the *only* things that matter for welfare, virtually all plausible theories of welfare hold that pains and pleasures are *relevant* to welfare, either directly or indirectly. Thus, perhaps, experiential features deserve more weight in the framework than, say, agential features⁴⁰ that may be relevant to fewer theories.

For simplicity I'm here assuming perfect theory-neutrality among a small number of theories. In reality, although we might want to consider multiple theories, we might lean toward some theories more than others. In that case, each feature would need to be weighted not only by the number of theories to which the feature is relevant but also by the plausibility of the theories to which the theory is relevant. And even if a feature is relevant to multiple theories, it may not be equally important to each theory. So the features would need to be weighted not only according to how many theories to which they are relevant, but also by how important they are to the theories for which they are relevant.

The final reason is that we may want to operationalize characteristics in more than one way. For instance, we may want to include a collection of physiological features⁴¹ in the database. Physiological features plausibly aren't intrinsically valuable; they're merely proxies for other categories. Suppose, for instance, we thought

general intelligence (whatever that means exactly) is pretty important for moral status or capacity for welfare. We might want to estimate general intelligence through a combination of neurological features (e.g., encephalization quotient, cortical neuron count) and behavioral features (e.g., tool use, uncertainty monitoring). Since no proxy will be perfect and it will be unclear which proxy is best, we will probably want to operationalize most characteristics in multiple ways. This needs to be reflected in the weighting system. If some characteristics are operationalized in more ways than others, then weighting features equally would amount to double-counting some characteristics.

Weighting the features looks important and inevitable. But doing so will be incredibly tough. Jean Kazez explains the difficulty thusly: “There are many capacities to which we assign positive value, but we don’t always have a definite idea of their relative values. If we’re trying to rank bower birds, crows, and wolves, it depends what’s more valuable, artistic ability (which favors the bower bird) or sheer intelligence (which favors the crow) or sociability (which favors the wolf). We’re not going to be able to put these three species on separate rungs of a ladder, in any particular order, and neither is the situation quite as crisp as a straightforward tie. We just don’t know how to assign them a place on the ladder, relative to each other” ([Kazez 2010](#): 87-88).

The trouble gets even worse when we consider so-called *combination effects*: “A property might raise the moral status of one being but not another, because it might raise moral status only when combined with certain other properties” ([Harman 2003](#): 177-178). For example, it might be the case that a certain degree of autonomy is required before some prosocial capacities contribute to moral status. Maybe nurturing behavior that is entirely pre-programmed and instinctive counts for less than love freely given. Honey bees and cows both care for their young, but if we think cows have a greater capacity for rational choice than honey bees, then the same level of juvenile guardianship might raise the moral status of cows more than honey bees.⁴² Thus, the feature weights may not be static or independent. Instead, they might be dynamic and interdependent. That is, the weight of a given feature may depend on the presence or absence of other features. Accounting for this complexity appears staggeringly hard.

In sum, accurate point estimates of the relative weights of the features are probably unachievable. The intellectually honest thing

to do is to assign each feature a range of weights.⁴³ But these ranges might be so wide as to rob the project of any action-guiding conclusions. That is, one’s views about the comparative moral value of different animals might depend almost wholly on how one interprets the relative importance of different features. If there’s nothing to be done to justifiably narrow the range of plausible weights, then the project may well end up countenancing a huge range of tradeoffs among species. If the project doesn’t tell us anything practical about which tradeoffs are permissible, then it is almost useless.

Conclusion

One cannot work in effective animal advocacy and wholly ignore the question of comparative moral value. Resources are finite, so tradeoffs among different animals are inevitable. Every time an organization launches a campaign, a researcher investigates a question, or a grantmaker funds a project, tradeoffs are made. The time, money, and attention that are devoted to one species could also have been devoted to another.

Although practical concerns will always guide our decision-making to some extent, it’s important to think about how we would *ideally like* to distribute resources. If the ideal distribution of resources differs significantly from what is currently feasible, we ought to devote time and money to surmounting those obstacles. It is not an immutable fact that, say, fish elicit little sympathy from the general public or that fish welfare organizations are few in number. If fish welfare is neglected relative to its importance, we can use our resources to improve tractability and overcome limiting factors.

To understand the ideal distribution of resources, we must understand the comparative value of different types of animals. If animal experiences and interests are all equal, then we should aim for interventions that have the greatest impact for the greatest number. Such a view would almost certainly push us to care more for small, numerous invertebrates than we currently do.

However, if some animals lack moral standing, then we should exclude those animals from our moral consideration. If that’s the case, it’s important to know where to draw the line between animals that have moral standing and those that do not. It’s also

42 In a recent [talk at Notre Dame](#), Eric Schwitzgebel offers the example of “a superpleasure machine but one with little or no capacity for rational thought. It’s like one giant, irrational orgasm all day long. Would it be great to make such things and terrible to destroy them, or is such irrational pleasure not really something worth much in the moral calculus?” Schwitzgebel is here wondering whether degree of rationality affects the moral value of capacity for pleasure, which would be another example of a combination effect.

43 It would also be helpful to assign probability distributions to the question ‘What weight would you assign to the feature after one hundred more hours of research?’ See [this comment from Nuno Sempere](#) about doing so in the context of investigating sentience.

important to know the consequences of drawing the line in one place rather than another, how confident we are in drawing the line, and what information would cause us to change where we draw the line.

If some animals have moral standing and others lack it, then by definition animals differ in their moral status and capacity for welfare. But there are plausible reasons to think that among the animals that have moral standing, there will be further variation in moral status or capacity for welfare. If that's the case, then saving the lives of one type of animal will not always have the same intrinsic moral value as saving the lives of other types of animals.⁴⁴ The suffering of, say, chimpanzees and octopuses may count for more, morally, than the suffering of, say, mealworms and prawn.

To determine the ideal allocation of resources, we need some way to measure capacity for welfare and moral status. Doing so will not be easy. The philosophical terrain is treacherous. Our intuitions are imprecise and likely skewed. The scientific literature is vast but uncertain. Still, even a small reduction in our uncertainty could make a big difference to our allocative decision-making. There are many years ahead to refine such estimates. But we can't put off making tradeoffs now.

Credits

This essay is a project of [Rethink Priorities](#). It was written by Jason Schukraft. Thanks to Kim Cuddington, Marcus A. Davis, Neil Dullaghan, David Moss, and Jeff Sebo for helpful feedback. If you like our work, please consider [subscribing to our newsletter](#). You can see all our work to date [here](#).

Works Cited

Anderson, D. J., & Adolphs, R. (2014). [A framework for studying emotions across species](#). *Cell*, 157(1), 187-200.

Bradley, A., Mennie, N., Bibby, P. A., & Cassaday, H. J. (2020). [Some animals are more equal than others: Validation of a new scale to measure how attitudes to animals depend on species and human purpose of use](#). *PLoS one*, 15(1), e0227948.

Browning, H. (2020). [Assessing Measures of Animal Welfare](#). *PhilSci Archive*, accessed May 2020.

Birch, J. (2017). [Animal sentience and the precautionary principle](#). *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 2(16), 1.

Buckwalter, W., & Stich, S. (2014). [Gender and philosophical intuition](#). *Experimental Philosophy Volume 2*, 307-346. Oxford University Press.

Budolfson, M., & Spears, D. (2019). [Quantifying Animal Well-Being and Overcoming the Challenge of Interspecies Comparisons](#). In B. Fischer (ed) *The Routledge Handbook of Animal Ethics* (pp. 92-101). Routledge.

Cauchoux, M., Chow, P. K. Y., Van Horik, J. O., Atance, C. M., Barbeau, E. J., Barragan-Jason, G., ... & Cauchard, L. (2018). [The repeatability of cognitive performance: a meta-analysis](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1756), 20170281.

Costa, A., Foucart, A., Hayakawa, S., Aparici, M., Apesteguia, J., Heafner, J., & Keysar, B. (2014). [Your morals depend on language](#). *PLoS one*, 9(4), e94842.

Cuddington, K., Fortin, M. J., Gerber, L. R., Hastings, A., Liebhold, A., O'Connor, M., & Ray, C. (2013). [Process-based models are required to manage ecological systems in a changing world](#). *Ecosphere*, 4(2), 1-12.

De Cruz, H. (2015). [Where philosophical intuitions come from](#). *Australasian Journal of Philosophy*, 93(2), 233-249.

Dicke, U., & Roth, G. (2016). [Neuronal factors determining high intelligence](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1685), 20150180.

Farrar, B. G., & Ostojic, L. (2019). [The illusion of science in comparative cognition](#). *PsyArXiv*. October 2, 2019.

Harman, E. (2003). [The potentiality problem](#). *Philosophical Studies*, 114(1), 173-198.

Herzog, H. (2010). The Importance of Being Cute. In [Some We Love, Some We Hate, Some We Eat](#), Harper Collins, 37-66.

Herzog, H., Grayson, S., & McCord, D. (2015). [Brief measures of the animal attitude scale](#). *Anthrozoös*, 28(1), 145-152.

44 Obviously, when comparing lives saved across species, differences in lifespan will need to be accounted for.

- Ioannidis, J. P. (2005). [Why most published research findings are false](#). *PLoS med*, 2(8), e124.
- Jardim-Messeder, D., Lambert, K., Noctor, S., Pestana, F. M., de Castro Leal, M. E., Bertelsen, M. F., ... & Herculano-Houzel, S. (2017). [Dogs have the most neurons, though not the largest brain: trade-off between body mass and number of neurons in the cerebral cortex of large carnivorous species](#). *Frontiers in Neuroanatomy*, 11, 118.
- Kagan, S. (2019). [How to count animals, more or less](#). Oxford University Press.
- Kahneman, D. (2011). [Thinking, fast and slow](#). Macmillan.
- Kazez, J. (2010). [Animalkind: What We Owe to Animals](#). Wiley-Blackwell.
- Kiyokawa, Y., & Hennessy, M. B. (2018). [Comparative studies of social buffering: A consideration of approaches, terminology, and pitfalls](#). *Neuroscience & Biobehavioral Reviews*, 86, 131-141.
- Machery, E., Mallon, R., Nichols, S., & Stich, S. P. (2004). [Semantics, cross-cultural style](#). *Cognition*, 92(3), B1-B12.
- MacLean, E. L., Hare, B., Nunn, C. L., Addessi, E., Amici, F., Anderson, R. C., ... & Boogert, N. J. (2014). [The evolution of self-control](#). *Proceedings of the National Academy of Sciences*, 111(20), E2140-E2148.
- McShea, D. W. (2017). [Logic, passion and the problem of convergence](#). *Interface Focus*, 7(3), 20160122.
- Miralles, A., Raymond, M., & Lecointre, G. (2019). [Empathy and compassion toward other species decrease with evolutionary divergence time](#). *Scientific Reports*, 9(1), 1-8.
- Norwood, F. B., & Lusk, J. L. (2011). [Compassion, by the pound: the economics of farm animal welfare](#). Oxford University Press.
- Rogers, L., & Kaplan, G. (2004). [All animals are not equal: the interface between scientific knowledge and legislation for animal rights](#). in C. R. Sunstein and M. Nussbaum (eds) *Animal Rights: Current Debates and New Directions*. Oxford: Oxford University Press, pp. 175-202.
- Schwitzgebel, E., & Cushman, F. (2012). [Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers](#). *Mind & Language*, 27(2), 135-153.
- Sebo, J. (2020). [How to Count Animals, more or less](#), by Shelly Kagan. *Mind*.
- Serpell, J. A. (2004). [Factors influencing human attitudes to animals and their welfare](#). *Animal Welfare* 13, S145-S152.
- Singer, P. (2011). [Practical ethics, 3rd Edition](#). Cambridge University Press.
- Swain, S., Alexander, J., & Weinberg, J. M. (2008). [The instability of philosophical intuitions: Running hot and cold on truth](#). *Philosophy and Phenomenological Research*, 76(1), 138-155.
- Walters, E. T., & Williams, A. C. D. C. (2019). [Evolution of mechanisms and behaviour important for pain](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 20190275.
- Weathers, S. T., Caviola, L., Scherer, L., Pfister, S., Fischer, B., Bump, J. B., & Jaacks, L. M. (2020). [Quantifying the Valuation of Animal Welfare Among Americans](#). *Journal of Agricultural and Environmental Ethics*, 1-22.
- Weiss, M. N., Franks, D. W., Croft, D. P., & Whitehead, H. (2019). [Measuring the complexity of social associations using mixture models](#). *Behavioral Ecology and Sociobiology*, 73(1), 8.
- Wuensch, K. L., Jenkins, K. W., & Poteat, G. M. (2002). [Misanthropy, idealism and attitudes towards animals](#). *Anthrozoös*, 15(2), 139-149.
- Wynne, C. D. L. (2007). [What are Animals? Why Anthropomorphism is Still Not a Scientific Approach to Behavior](#). *Comparative Cognition & Behavior Reviews*, 2, 125-135.