

How to Express Improvements in Animal Welfare in DALYs-averted

BOB FISCHER*

TABLE OF CONTENTS

I. INTRODUCTION	735
II. CLARIFYING THE SIGNIFICANCE OF THE CONCEPTUAL PROBLEM . . .	736
III. BASIC STRATEGIES FOR ADDRESSING THE CONCEPTUAL PROBLEM	739
IV. A SOLUTION TO THE CONCEPTUAL PROBLEM	741
V. ESTIMATING WELFARE RANGES	743

I. INTRODUCTION

Effective altruists (EAs) want to do as much good as possible with the charitable resources available to them. If EAs want to do the most good per dollar spent—that is, if they want to maximize cost-effectiveness—they need a common currency for comparing very different interventions. GiveWell, Open Philanthropy, Founders Pledge, and many other effective altruist organizations currently have such a metric: namely, the number of “disability-adjusted life years (DALYs) averted.” A DALY is a health measure with two parts: *years of life lost* and *years lost to disability*. The former measures the extent to which a condition shortens a person’s life; the latter measures the health impact of living with a condition in terms of years of life lost. Together, these values represent the overall burden of the condition. So, *averting* a DALY is averting a loss—namely, the loss of a single year of life that is lived at full health.

There are several standard concerns about DALYs, both as a health measure and as a tool for policymaking.¹ I set those concerns aside here. Instead, I will focus on a well-known implication of wanting to do the most good *impartially understood*: namely that if it is possible to do good by helping nonhumans, then it

* Bob Fischer is an Associate Professor of Philosophy at Texas State University, a Senior Research Manager for Rethink Priorities, and the Director of the Society for the Study of Ethics and Animals. © 2023, Bob Fischer.

1. Sudhir Anand & Kara Hanson, *Disability-adjusted Life Years: A Critical Review*, 16 J. HEALTH ECON. 685, 685–702 (1997); Trude Arnesen & Erik Nord, *The Value of DALY Life: Problems with Ethics and Validity of Disability Adjusted Life Years*, 319 BRIT. MED. J. 1423, 1423–25 (1999); Erik Nord, *Disability Weights in the Global Burden of Disease 2010: Unclear Meaning and Overstatement of International Agreement*, 111 HEALTH POL’Y 1, 99 (2013).

is essential to be able to compare ways of doing good that benefit humans and nonhumans alike. But to assess how much good they can do by helping animals, EAs need to convert benefits to animals into whatever terms they use to assess other causes. If DALYs-averted is the relevant unit of goodness, then that means converting benefits to animals into DALYs-averted. In other words, EAs need to express those benefits in terms of a year of additional human life that is lived at full health.

But there's a conceptual problem here. The DALY is a health measure, not a welfare measure, and the two are not perfectly correlated: you can be in good health and have poor welfare (ask anyone who has grieved a loss without being clinically depressed); likewise, you can be in poor physical health and have good welfare.² However, animal welfare science—our main source of evidence about what is good and bad for animals—attempts to assess welfare, not health. While animal welfare scientists disagree about the best theory of animal welfare, there is considerable momentum behind some kind of hedonism: things go well for animals insofar as they have positively valenced conscious states; things go poorly for animals insofar as they have negatively valenced conscious states. So, health economists have health measures, animal welfare scientists have (hedonic) welfare measures, and these tools are not designed to play together.

My aim in this paper is to develop one solution to this conceptual problem. The paper proceeds as follows. In Section II, I explain the significance of the conceptual problem in more detail. In Section III, I sketch some potential solutions to the problem and explain why they are not satisfying. In Section IV, I lay out my own proposal.

II. CLARIFYING THE SIGNIFICANCE OF THE CONCEPTUAL PROBLEM

The conceptual problem stems from the sources of evidence on which effective altruists rely: global health literature (when focused on humans) and animal welfare science (when focused on animals). If we want to understand the conceptual problem more fully, it may help to try to dismiss it. So, let's consider the charge that we're dealing with a pseudo-problem. We can imagine the critic arguing as follows:

True enough, EAs want to do the most good per dollar spent. True enough, the best sources of information for assessing whether they're doing the most good in the near term—health economics literature and animal welfare science literature—use measures that weren't designed to be directly comparable. But even if the DALY isn't a direct *measure* of welfare, it is clearly a *proxy* for welfare. And while animal welfare scientists may be interested in measuring hedonic welfare, they do not know how to measure it directly. In practice, then, they're often measuring indicators of health. So, while these measures

2. See Sonja Lyubomirsky, *Hedonic Adaptation to Positive and Negative Experiences*, in THE OXFORD HANDBOOK OF STRESS, HEALTH, AND COPING 200–24 (S. Folkman ed., 2012).

are different in *theory*, they may be fairly similar in *practice*. That is, while health measures are not welfare measures, the two might be correlated well enough for many purposes, cost-effectiveness analyses among them.

There are two problems with this objection. First, even if it is true that the two measures are sufficiently correlated for some practical purposes, this objection concedes that they are not the same. Therefore, insofar as there's any value in getting our conceptual house in order, there is a project to pursue.

Second, the spirit of this objection is that the conceptual problem does not matter because it is not a *practical* problem. But that is wrong—or, at least, many will think so. Suppose we blur the distinction between the DALY framework and the welfare assessment frameworks that animal welfare scientists employ, and directly express an assessment of, say, layer hen welfare as a DALY burden. Consider, then, the lives of layer hens in conventional cages (versus, say, an aviary system), which are often thought to be quite poor.³ In the DALY framework, health conditions receive scores between zero and one, with one being the worst. In humans, for instance, a fractured femur is ~ 0.3 . Since most layer hens in conventional cages have keel bone fractures, we might say that the DALY burden for life in a conventional cage should be at least 0.3.⁴ But then we would be saying that a chicken spending a year in a conventional cage is at least as bad as the loss of roughly four months of human life lived at full health. And if *that's* right, then given how many layer hens there are, there's a good chance EAs should be spending a lot more money trying to help them.

This is not news: EAs are well aware of the issue. Holden Karnofsky, co-CEO of Open Philanthropy, has commented on just this issue:

Some people think that animals such as chickens have essentially no moral significance compared to that of humans; others think that they should be considered comparably important, or at least 1-10% as important. If you accept the latter view, farm animal welfare looks like an extraordinarily outstanding cause, potentially to the point of dominating other options: billions of chickens are treated incredibly cruelly each year on factory farms, and we estimate that corporate campaigns can spare over 200 hens from cage confinement for each dollar spent. But if you accept the former view, this work is arguably a poor use of money.⁵

Essentially, blurring the distinction between the DALY framework and animal welfare assessment is tantamount to saying that chickens (among other animals)

3. For a description of the conditions, see BOB FISCHER, *ANIMAL ETHICS: A CONTEMPORARY INTRODUCTION* (2021). For quantitative assessments of these animals' welfare, see F. BAILEY NORWOOD & JASON L. LUSK, *COMPASSION, BY THE POUND: THE ECONOMICS OF FARM ANIMAL WELFARE* (2011).

4. Christina Rufener & Maja M. Makagon, *Keel Bone Fractures in Laying Hens: A Systematic Review of Prevalence Across Age, Housing Systems, and Strains*, 98 *J. ANIMAL SCI.* 1, S36–S51 (2020).

5. Holden Karnofsky, *Worldview Diversification*, OPEN PHILANTHROPY (Dec. 13, 2016), <https://www.openphilanthropy.org/research/worldview-diversification/> [<https://perma.cc/KLQ8-EAUI>].

are as important as humans. Given their numbers and the harms imposed on them, campaigns to help are probably among the most cost-effective interventions available. Indeed, Karnofsky goes on to report that based on Open Philanthropy’s own estimates, “if you value chicken life-years equally to human life-years . . . [then] corporate campaigns do about 10,000x as much good per dollar as top [global health] charities.”⁶ Two more recent estimates agree that corporate campaigns are more cost-effective. If we assign equal weight to human and chicken welfare in the model that uses, corporate campaigns are roughly 5,000 times better than the best global health charities.⁷ If we do the same thing in the model that Clare and Goth employ, corporate campaigns are 30,000 to 45,000 times better.⁸

Open Philanthropy handles uncertainty about the relative moral significance of chickens via “worldview diversification,” which, for practical purposes, means having several siloed budgets. The relevant ones here are a budget devoted to animals and a budget devoted to more traditional public health and development programs, such as deworming campaigns, distributing bed nets, and cash transfers to the world’s poorest people. But, of course, Open Philanthropy has to decide the relative sizes of those budgets. Currently, animals receive a relatively small portion of neartermist dollars (i.e., those devoted to animals and traditional human-oriented global health and development causes, in contrast with Open Philanthropy’s longtermist spending, which I ignore here). If we blur the distinction between welfare assessment and the DALY framework, then due to the sheer number of layer hens who can be helped—just to mention one of many animals, and certainly not the most numerous—it would be surprising if that arrangement were optimal.

Here is the upshot: even if there is some sense in which health economics and animal welfare science are measuring closely related things—as the initial objection maintains—there is clearly some *other* factor (or set of factors) that many EAs think these measurements are missing. These EAs do *not* think that the good represented by averting the loss of a year of human life at full health is equivalent to the good represented by averting the loss of a year of chicken life at full health. So, if a straightforward conversion of animal welfare assessments into the DALY framework suggests otherwise—where, again, we assign scores between zero and one to animals’ health conditions based on parallels with humans’ health

6. *Id.*

7. Vasco Grilo, *Corporate Campaigns For Chicken Welfare are 10,000 Times as Effective as GiveWell’s Maximum Impact Fund?*, EFFECTIVE ALTRUISM F. (July 28, 2022), <https://forum.effectivealtruism.org/posts/nDgCKwJBKwFvcBsts/corporate-campaigns-for-chicken-welfare-are-10-000-times-as-https://perma.cc/F8ZM-CJQ3>.

8. Stephen Clare & Aidan Goth, *How Good is The Humane League Compared to the Against Malaria Foundation?*, EFFECTIVE ALTRUISM F. (Apr. 29, 2020), <https://forum.effectivealtruism.org/posts/ahr8k42ZMTvTmTdwM/how-good-is-the-humane-league-compared-to-the-against/> [https://perma.cc/556H-85F9]. This range emerges from the default set of parameters in their respective Guesstimate model after setting the node “moral weight (DALY/cDALY)” to one. It is a range because their Guesstimate model is noisy and different samples give different results.

conditions—there is something wrong with such a conversion. For these EAs, anyway, the conceptual problem is real.

III. BASIC STRATEGIES FOR ADDRESSING THE CONCEPTUAL PROBLEM

That being said, it is unclear what exactly the straightforward conversion misses. Let's consider some candidates.

One radical possibility is that the relevant animals are not even moral patients. For instance, perhaps these animals are not sentient. If sentience is necessary and sufficient for moral status, then these animals do not count morally at all.

Granted, there are indeed theories of consciousness—such as certain higher-order theories—that imply that many nonhuman animals are not conscious and, therefore, not sentient. However, there is considerable uncertainty about the correct theory of consciousness and, despite this, a general consensus among consciousness scientists that many farm animals *are* sentient. So, while it might be reasonable to believe a given theory (or family of theories) that implies that, say, chickens are not sentient, this seems like a reason to *discount* chickens' welfare based on uncertainty about their being moral patients—not a reason to deny it entirely. Even if we assign a relatively low credence to chickens being sentient—say, 0.6, where the hypothesis is only somewhat more likely than not—this would hardly affect the practical conclusion once we multiply through. To prevent chickens from dominating, we would need something that drops their significance by orders of magnitude, not by less than a factor of two.

Alternately, then, someone might insist that the straightforward conversion from animal welfare assessments to DALYs overlooks the possibility that *hierarchicalism* is true, according to which equal quantities of welfare are not necessarily equally morally important. On this view, some individuals' welfare, or some kinds of welfare, matter more than others—a view now associated with Shelly Kagan.⁹ The alternative is unitarianism, according to which all welfare matters equally—a view usually associated with Peter Singer.¹⁰ The main challenge for hierarchicalists is to tell us why some individuals' welfare matters more than others', which is no small feat. Without such an explanation, hierarchicalism seems to be objectionably arbitrary. Moreover, insofar as hierarchicalists lack such an explanation and still insist that degrees of mattering track species boundaries, they are vulnerable to the charge of speciesism.

Why is it difficult to explain why some individuals' welfare matters more than others'? There are two main problems. First, suppose we appeal to some trait to explain why an individual's welfare matters more than another's. Then we will always face this question: why does that trait make an individual's welfare matter more? Why doesn't it just show that they have more welfare at stake in many situations? Without an answer, it looks like hierarchicalists are double-counting differences between humans and animals: they use differences in traits to argue for

9. SHELLY KAGAN, *HOW TO COUNT ANIMALS, MORE OR LESS* (2019).

10. PETER SINGER, *ANIMAL LIBERATION: A NEW ETHICS FOR OUR TREATMENT OF ANIMALS* (1975).

differences in interests, which then justify differential treatment. Then, they also use differences in traits to argue that even when humans and animals have similar interests, humans' interests matter more. But if the interests really are similar, that seems like speciesism.

Second, to accommodate widely-held intuitions, many people want there to be differences between whole groups of individuals—usually species. For instance, they want humans to matter more than chickens, who matter more than carp, who matter more than shrimp, and so on. But at every level in this hierarchy, they have to appeal to traits that explain the difference. And at every level, (a) they are bound to find some members of the group who lack the relevant trait and (b) they are bound to find some members of other (lower) groups who have it. Hence, either their hierarchy will not track species boundaries, due to significant intraspecies variation, or it will be arbitrary. Most people find the former unacceptable, so yet again the problem of arbitrariness looms large.

Moreover, even if we get solutions to these problems, it is not clear that hierarchicalism helps unless it is overtly speciesist—i.e., unless it assumes that species membership *per se* is the morally relevant factor that explains why humans have a higher moral status than nonhuman animals, rather than a proxy for some morally relevant factor(s). This is because it seems unlikely that any hierarchicalist position is going to block the conclusion that EAs ought to be spending more neartermist dollars on nonhuman animals. To see this, just consider how we might assess what a reasonable discount rate might be for various individuals' welfare or kinds of welfare. Presumably, we would (1) identify the various bases for discounting some welfare relative to others; (2) where possible, identify *humans* whose welfare would be discounted on each basis; and (3) where no human's welfare would be discounted on the relevant basis, consider sympathetic animals whose welfare would be discounted on that basis (e.g., cats and dogs). Once we do that, we are likely to find that peoples' intuitions support relatively modest discount rates—if any at all—effectively gutting the practical significance of hierarchicalism.

This leaves an option that many philosophers have suggested: that humans and nonhumans generally realize different amounts of welfare.¹¹ If this view is correct, then there is an obvious problem with the straightforward conversion of animal welfare assessments into the DALY framework. Essentially, it assumes that when a chicken has a certain health condition, like a keel bone fracture, that

11. See, e.g., Mark Budolfson & Dean Spears, *Public Policy, Consequentialism, the Environment, and Non-human Animals*, in THE OXFORD HANDBOOK OF CONSEQUENTIALISM 592 (D. W. Portmore ed., 2020); KAGAN, *supra* note 9; JEFF MCMAHAN, THE ETHICS OF KILLING: PROBLEMS AT THE MARGINS OF LIFE (2002); JEFF SEBO, SAVING ANIMALS, SAVING OURSELVES: WHY ANIMALS MATTER FOR PANDEMICS, CLIMATE CHANGE, AND OTHER CATASTROPHES (2022); SINGER, *supra* note 10; Peter Vallentyne, *Of Mice and Men: Equality and Animals*, in EGALITARIANISM: NEW ESSAYS ON THE NATURE AND VALUE OF EQUALITY 211–238 (N. Holtug & K. Lippert-Rasmussen eds., 2007); Kevin A. Wong, Counting Animals: On Effective Altruism and the Prospect of Interspecies Commensurability (Dec. 2016) (B.A. thesis, Princeton University) (on file with DataSpace at Princeton University).

chicken is roughly as badly off as a human with an ostensibly similar health condition, such as a broken leg. But if animals usually do not realize as much positive or negative welfare as humans, then we cannot make this assumption. While we can still score animals' health conditions on a zero to one scale, it might work out that keel bone fractures deserve a score of 0.00003—not 0.3, like a fractured femur in a human.

IV. A SOLUTION TO THE CONCEPTUAL PROBLEM

Suppose we accept that humans and nonhumans can realize different amounts of welfare. This is not yet a solution to the conceptual problem, as it is not yet precise about the relationships between welfare, health conditions, and the relevant welfare and health assessment tools. But we are close. Here, I think, is the way to close the gap.

The basic proposal involves *welfare ranges*. Every individual has a certain capacity for welfare—i.e., how much welfare they *could* realize in total (even if, in fact, they actually realize much less). An individual's capacity for welfare has two components: their welfare range (which is just their capacity for welfare at a given time) and their lifespan. Since lifespans are relatively easy to estimate—and are often irrelevant because we're looking at welfare impacts over fixed periods—we will ignore them here.

An individual's welfare range is the difference between the best and worst welfare states the individual can realize at a time. In other words, assume we can assign a positive number to the best welfare state the individual can realize and a negative number to the worst welfare state the individual can realize. The difference between them is their welfare range. If we define welfare ranges in a certain way, and estimate their size, then we can estimate how much welfare humans realize relative to various nonhuman animals.

Ultimately, the goal is to convert welfare level changes—i.e., differences in *realized* welfare, not welfare ranges—into DALYs-averted. So, while an individual's welfare range is the difference between the “best” and “worst” welfare states that the individual can realize, we are not interested in the best and worst *in principle*. Instead, we are interested in the interpretations of “best” and “worst” that lend themselves to converting welfare level changes into the relevant unit of goodness. Given that DALYs-averted is the relevant unit of goodness, I propose we take a DALY-averted as a proxy for some amount of welfare: namely, the amount of welfare that the average human in full health realizes over the course of a year. Then, we can say that the best *human* welfare state is the average welfare level of the average human in full health.

Plainly, there are many other senses of “best” where humans realize more welfare still. Full health is hardly the pinnacle of welfare: for instance, there's being at full health while having a peak experience of some kind (e.g., holding your newborn child), which is no doubt superior to being at full health *simpliciter*. However, we are not *excluding* those welfare levels. They are included in the average. We just do not want to focus on them. If we did, then we would risk

overstating how much welfare a DALY-averted represents. It is not equivalent to the *maximum* amount of welfare that humans-in-full-health *could* realize over the course of a year, but instead, to some lower amount, though not one we can specify. However, as we will see below, we do not need to be able to specify the amount. We just need to index our welfare range estimates to it. Accordingly, the relevant best *animal* welfare states will be analogous: the average welfare level of the average animal (of a given species) in full health.¹²

It is important to note that the DALY framework does not allow for states worse than death. Effectively, it assumes that humans' net welfare range is exclusively positive. This is not to say that humans cannot have any negative welfare states, just that their welfare is net positive even when experiencing those states. There are probably good moral reasons to bake this assumption into the framework, given inevitable disagreement about which lives are worth living and the DALY framework's use in resource allocation. However, whatever the merits of the net positivity assumption, it is controversial as a theoretical claim. Still, it will not make any difference here whether we accept it; therefore, we can take it on for simplicity's sake. However, there is no parallel moral reason for assuming that many nonhuman animals have net positive lives. Indeed, it is quite plausible that many chickens, for instance, have net negative welfare states in standard intensive farming systems. So, let's *not* assume that their lives are invariably net positive.

Finally, if we allow that animals can be in states worse than death, we face the question of the *skews* of their welfare ranges. Can they be worse off than they can be well off? Better off than they can be badly off? Or are their welfare ranges symmetrical around the neutral point, such that they can be equally well off and badly off? Though this topic deserves considerable attention, it would take us too far afield here. Thus, we will assume symmetry: if the "best" welfare state for a chicken is represented by some arbitrary positive number, then the "worst" welfare state is represented by the negation of that number.¹³

Now suppose that we can estimate differences in welfare ranges. The next step is to reinterpret the welfare assessments that come out of animal welfare science. We can understand animal welfare assessments as *species-relative* welfare assessments, measures of how well or poorly things are going for nonhuman animals relative to what's best and worst *for them*. They are not estimates of the absolute amount of welfare that nonhuman animals realize relative to humans in any given case. Given this, we should convert the results of welfare assessments into a common unit.

12. We may need to add some caveat about living in reasonably hospitable circumstances, as we'd expect of a person living in full health.

13. Assuming symmetry may stack the deck against animals, as it is plausible, they can fare worse than they can fare well. But we will ignore that complication here, as anything that stacks the deck against animals cannot be criticized by those committed to large difference in welfare ranges between humans and nonhumans.

Suppose, for example, that chickens' welfare range is 10 percent of humans' welfare range. Additionally, suppose that over the course of a year, the average chicken is about half as badly off as they could be in a conventional cage (i.e., their average welfare level welfare is roughly halfway between their neutral point and their worst welfare state) and a quarter as badly off as they could be in an aviary system (i.e., their average welfare level is roughly a quarter of the way from their neutral point toward their worst welfare state). Then we can calculate the equivalent amount of human welfare and, by extension, the number of DALY-equivalents averted:

1. Chickens' welfare range is 10 percent of human's welfare range.
2. Assuming symmetry around the neutral point, the negative portion of chickens' welfare range is 5 percent of human's positive welfare range (e.g., if humans' welfare range is 100 and chickens' welfare range is ten, chickens range from negative five to five, and the negative portion of that is 5 percent of humans' welfare range).
3. Given our assumptions about the welfare impacts of the two production systems, the move from conventional cages to aviary systems averts an amount of welfare equivalent to 25 percent of the average chicken's negative welfare range (e.g., and continuing with the numbers in the previous example, it moves chickens from -2.5 to -1.25).
4. So, assuming symmetry, 25 percent of chickens' negative welfare range is equivalent to 1.25 percent (5 percent \times 25 percent) of humans' positive welfare range.
5. By definition, averting a DALY averts the loss of an amount of welfare equivalent to the positive portion of humans' welfare range for a year.
6. So, assuming symmetry, the move from conventional cages to aviary systems averts the equivalent of 0.0125 DALYs per chicken per year on average.

That, in a nutshell, is my solution to the conceptual problem with which we began. It solves the problem by treating a DALY-averted as representing some quantity of welfare, defining welfare ranges in terms of DALYs-averted, and accepting a species-relative interpretation of animal welfare assessments.

Of course, the practical upshot of the solution depends on the relative welfare range estimate. I have not said anything about how we might assess relative welfare ranges. But that is not the burden here. The goal is just to show how to convert units from one evidential source (animal welfare science) into units from another (health economics). And that is what I have done.

V. ESTIMATING WELFARE RANGES

Before I conclude, briefly consider the task of estimating differences in welfare ranges across species.¹⁴ The research program is based on the idea that animals

14. Individuals have welfare ranges; species do not. Still, we can understand talk about a species' welfare range as referring to the average welfare range of that species' members.

may vary with respect to their ability to realize the determinants of welfare. Some animals may have more intense pains than others; some may be able to acquire more knowledge; others may make deeper, more significant relationships; others still may have richer, more complex desires. So, given a theory of welfare, we can identify the determinants of welfare. With an understanding of the determinants of welfare, we might be able to find empirical evidence that reflects relevant variation in the ability to realize the determinants of welfare. In other work, I have spent considerable effort on exactly this task.¹⁵

Of course, there are many sources of uncertainty along the way. There is uncertainty about both the correct theory of welfare and how best to handle axiological uncertainty. This would not be a problem if theories of welfare were largely in agreement about welfare ranges, but they do not appear to be. Consider a garden variety objective list theory according to which the following goods contribute positively to welfare: acting autonomously, gaining knowledge, having friends, being in a loving relationship, doing meaningful work, creating valuable institutions, and so on. While some might object to the application of this theory of welfare to nonhuman animals,¹⁶ it remains the case that if we do apply it to them¹⁷ many nonhuman animals will have smaller welfare ranges than many humans. Presumably, the empirical facts will indicate that they cannot realize many of the goods on the list. By contrast, consider a simple version of hedonism that rejects the higher/lower pleasure distinction. If we assume such a theory, we might expect much smaller differences in welfare ranges between many humans and animals.

Similarly, even given a theory of welfare, there is uncertainty about what constitutes evidence of variation in the ability to realize the determinants of welfare. Almost everyone agrees that positively and negatively valenced experiences—pleasures and pains—are welfare-relevant. But there are no good direct objective measures of relative pain intensity across humans, much less across species. Likewise, there are no good objective measures of the relative strengths of desires, the relative depths of romantic relationships, the meaningfulness of activities, or many other potential determinants of welfare.

Of course, the absence of direct objective measures necessitates the use of decent proxies for such differences, though the switch to indirect measures raises its own uncertainties. If we are using proxies for the variation in the capacity to generate the determinants of welfare, we have to figure out how to score them and aggregate the scores. For instance, it is hardly clear how many levels of pain—or desire strength, or friendship, or anything else—there may be.

The upshot: I will be the last to suggest that it is easy to estimate differences in welfare ranges. In principle, though, we can make progress. And given the stakes for resource allocation in EA, it is essential to try.

15. WEIGHING ANIMAL WELFARE: COMPARING WELL-BEING ACROSS SPECIES (Bob Fischer ed.) (forthcoming 2024).

16. Donald W. Bruckner, *Human and Animal Well-Being*, 102 PAC. PHIL. Q. 393, 393 (2021).

17. See Eden Lin, *Welfare Invariabilism**, 128 ETHICS 320, 331 (2018).