



Moral Weights of Animals, Considering Viewpoint Uncertainty

Richard Bruns¹ and Jim Davies²

Many utilitarians would like a number to use to evaluate the moral impact of actions that affect animals. However, there is a great disagreement among scholars involved with animal ethics, both about how much different animals can suffer and how much that suffering morally matters. To illustrate this uncertainty, while showing as a proof of concept that it may be possible to produce useful estimates in spite of it, we ran a Monte Carlo simulation that samples the ranges of major viewpoints scholars hold in the field, to show a spread of uncertainty for how we should treat six representative animals: crickets, salmon, chickens, pigs, cows, and elephants. The results show that the uncertainty is very large, with a 90% confidence interval ranging between an animal having no value and being valued as much as a human being. More research, in the form of expert surveys and a thorough and rigorous literature review, would be required to produce better estimates, but as an illustration, we present 20% and 40% confidence intervals, as well as the median and geometric mean, based on weighting the theories according to our informal estimate of their prevalence in the literature.

Introduction

Numerical estimates are important for use in utilitarian calculations that inform how we should treat animals (Sandhoefner 2017; Tomasik 2018). To determine these numbers, we must determine *how much* their welfare has changed and the *moral value* of that change.

For example, some scholars of the past have suggested that non-human animals (hereafter “animals”) felt nothing at all, and as such could not have welfare or any change to it. Some believe that animals feel things to a lesser extent than humans do. Others believe that animals feel things the same way that humans do, and a few suggest that, at least in some situations, they might feel more.

Even if this problem were decided, we still have disagreement on the moral implications of animal suffering: do animals have no moral status, the same moral status as humans, or less moral status than humans?

It is tempting for individuals to make up their own minds on these issues, and then conduct a utilitarian calculation based on those decisions. Our approach is different.

¹ Johns Hopkins Center for Health Security; corresponding author: richardbruns@gmail.com

² Carleton University



Rather than arguing for one particular view on these controversial issues, we treat experts' disagreement as uncertainty in the field, and try to describe the spread of this uncertainty (see MacAskill 2014). We demonstrate an approach for polling experts and aggregating their viewpoints numerically, while also accounting for scientific uncertainty in the inputs to their chosen theories.

In this paper, we model the uncertainty in the scholarly literature with a Monte Carlo simulation. During each iteration, it chooses one viewpoint from among the scholarly opinions, in addition to sampling from the input ranges of various relevant characteristics of animals. After running thousands of iterations, we can graph the results to visually see the spread of disagreement, and calculate the statistical properties of the results.

Our main contribution is not our results; it is the demonstration of a method for producing better results. A structured survey process would be required to first produce a definitive list of possible moral theories of suffering and value, and then to find the credence that experts assign to these theories. This survey process should then be conducted regularly, so we can update our numbers as the scholarly consensus changes through research and debate.

We focus on animals typically eaten (chickens, salmon, pigs, and cows) as well as small animals that some advocate we eat (crickets), and the African Elephant, as an example of an animal that is larger than humans. This list has the benefit of having examples from insects, birds, and mammals.

Theories of Suffering

This paper focuses on the extent to which animals feel a conscious change in their own welfare, which we will call "suffering." This isn't a perfect term, as some have defined suffering as a second-order attribution to a first-order sensation, one that, perhaps, requires a sense of self and/or a sense of time (Dennett 1995). However, the term "pain" is also problematic, as it does not require consciousness, and has been shown in humans to be dissociable from the "awfulness" of pain (that is, under certain analgesics, people claim to be aware of the pain but not mind it; Shriver 2009). By suffering, we mean that the creature consciously feels something unpleasant.

Given the same event, how much will a given kind of being suffer? Here we have four basic views: The "Unconscious Animal" view holds that though animals act as though they suffer, they have no consciousness and therefore do not. Though not a popular view nowadays, it was famously held by René Descartes (see Huebner 2011 for a discussion), but some contemporary scholars endorse this view as well (Harrison 1991). Some believe that ray-finned fish are conscious of nothing (Key 2015). Beliefs in the Unconscious Animal theory can be supported through some Higher-Order Thought theories of consciousness. On some interpretations, if animals are incapable of



higher-order thought, then they would similarly be incapable of having conscious mental states (Carruthers 2000, though his view has recently changed).

The “Muted Consciousness” view holds that animals suffer (and experience joy) less than humans do, given the same event (Scherer, Tomasik, Rueda and Pfister 2018).¹ Specifically, these scholars believe that the psychological complexity of the animal in question is thought to determine how muted its feelings are. So, for instance, a cow breaking a femur would suffer more than a frog breaking a femur. How best to measure the relevant psychological complexity is also a matter of debate. Research has been done on measuring the behaviors of various animals, but there is currently no way to map these measures of behavioral complexity onto a single scale. As an illustration, we use numbers that can be compared across animals, and might plausibly affect psychological complexity: brain mass, brain-body ratio, Encephalization Quotient, number of neurons, and number of cortical neurons. These different sub-theories result in different multipliers between 0 and 1.

We assume a linear mapping between these numbers and capacity for suffering, although future research may show that a different mapping should be used. Calculation of all of these biological traits (body mass, cortical neuron counts, and so on) requires empirical findings. The sources of these numbers are in the Appendix. When determining a suffering factor for animals, we use the animal number divided by the analogous number for humans. For example, a cow’s Encephalization Quotient is about 0.58, and a human’s is about 7.6, so according to this theory the cow feels about 0.076 ($0.58/7.6$) times, or about $1/13^{\text{th}}$, as much suffering as a human would from the equivalent event.

The third view, which we will call the “Same Suffering” view, is that, given the same event, all animals that can feel anything at all suffer with an equal intensity. Thus, a human, cow, or mouse breaking a femur results in an equal amount of well-being lost for each animal. For this theory, the multiplier would be 1.

Finally, we have the view that, in certain cases, at least, animals can feel *more* suffering for the same event. The reason for suggesting this is that humans have many interests, and being pain-free is just one of them. There is no name for this theory, as far as we know, so we will call it the “Tinker Bell Theory,” in reference to Barrie’s explanation that Tinker Bell’s emotions completely consume her because she is so small that she has no room for complex emotion.² Thomson (1990, 292) suggests that, all else being equal, it is worse to put animals in pain than humans for this reason. Akhtar (2011) argues that animals can feel more suffering from pain than humans, but acknowledges that human complexity can also make pains worse.

¹ Another name for this is the “unequal interests” model (DeGrazia, 2008).

² This idea is expressed in the stage directions of the play *Peter Pan*, in reference to Tinker Bell’s mental states: “She is not wholly heartless, but is so small that she has only room for one feeling at a time.”



The effects of age, visualization, and meditation on pain have been studied empirically, giving us a plausible order-of-magnitude bound of the effect size of humans' ability to use thinking to reduce the unpleasantness or intensity of pain.

It is possible that human infants are like animals, in that they have not yet developed cognitive strategies to help them deal with pain, and as such would suffer more from pain than adult humans would. One study compared the brain response of newborns and adults from being poked, and found that the same neural response required about 1/4 as much force in newborns (Goksan et al. 2015). We do not know if this is because of differences in skin thickness, transduction, or something else, but if it reflects differences in true pain sensitivity, we might use 4 as an upper-bound rough estimate for the sensitivity to pain of an animal.

We can also look at specific methods adults use to deal with pain and compare them to people who do not use those methods. Zeidan, Martucci, Kraft, Gordon, McHaffie and Coghill (2011) show that meditation reduced the unpleasantness of pain by 57%. A study of expert meditators found a similar reduction of about 59% in pain unpleasantness (Lutz, McFarlin, Perlman, Salomons, and Davidson 2013). This suggests that non-meditators feel double or triple the suffering from pain that meditators do.

If we assume that meditators are the top end of a 'mindfulness spectrum' and animals are at the bottom end, then it is possible that animals experience even more unpleasantness from pain than humans do. However, it is also possible that meditation turns off the narrative part of the human mind that allows the unique kind of suffering that humans can experience, and that animals with no narrative thoughts perceive pain like meditators do.

Looking at the suffering of newborns and meditators, we assign a uniform distribution between 1/3 and 4 for the Tinker Bell estimate, with very low confidence.

Theories of Value

In the previous section, we discussed different theories on how an event or situation would affect an organism in terms of causing suffering. In this section we will hold suffering constant and ask whether the same suffering in different kinds of beings results in a greater or lesser moral value.

There are three main theories we address. The first we will call "Extreme Speciesism." For this view, animals might be sentient, but they don't matter (or only matter instrumentally). Carruthers holds this view, because he bases moral status on the ability to enter into contracts, which animals cannot do (2011). Mathematically, this assigns a moral weight of 0 to animals. The second view is the "Equal Consideration" view, which holds that moral goodness and badness depends on welfare, but it doesn't matter what



creature is experiencing the welfare (DeGrazia 2008, 189).³ Third is the “Unequal Consideration” view.⁴ This view holds that even when the suffering is the same, we should give more moral consideration to humans than other animals (and, perhaps, some animals more than others.)

Choosing Weights

Crucial to this analysis are weights for these theories. Our approach is based on epistemic humility. We do not claim to know the truth, or that objective truth exists for moral questions.⁵ Therefore, we use weights based on how many scholars in relevant areas (comparative psychology, biology, animal ethics, cognitive science, etc.) endorse each of these four theories. When an iteration of the simulation picks a theory, it does so based on this weight. However, such a survey of experts has not been done, so we have guessed at the prevalence of these theories based on a general reading of the literature. We freely acknowledge that these weights can be seen as arbitrary, and that they are likely to be heavily influenced by what we consider to be reasonable. We use them to show that it is possible to aggregate both mainstream and fringe theories in a sensible way, rather than simply assigning equal weight to all (or a selection of) theories.

Our estimates for the weights of theories of suffering are as follows: Unconscious Animal Theory: 3%. Muted Animal Theory: 62%. Same Suffering Theory: 30%, Tinker Bell Theory: 5%. The Muted Animal Theory can further subdivide according to what measure of complexity people use to determine how much muting is going on. These are our estimates, which sum to the 62%: Raw Brain Mass: 2%, Brain-Body Ratio: 5%. Encephalization Quotient: 20%. Neuron Count: 15%. Cortical Neuron Count: 20%. These various complexity measures rely on empirical facts about different animals (brain mass, etc.).

Our estimates for theories of value are: Extreme Speciesism: 10%. Equal Consideration: 60%. Unequal Consideration: 30%.

Methods

We ran a Monte Carlo simulation with 10,000 iterations. Each iteration calculates a weight for each animal considered (cricket, salmon, chicken, pig, cow, elephant),

³ Also known as the “welfarist view” (Copp, 2011).

⁴ Shelly Kagan (2019) calls the “Equal Consideration” view the “unitarian” view, and the “Unequal Consideration” view the “hierarchical” view.

⁵ If you are a moral realist, then this approach tells you what the probability distribution of your belief should be, under the assumption that all of the scholars are trying to find truth and each one has the same probability of being correct. If you are a moral anti-realist, then this approach is useful to aggregate preferences to form a shared social decision about what to do regarding animals. We will note that one author strongly believes in moral realism, and the other strongly believes in moral anti-realism, and that we both find this approach to be helpful to us.



according to credence estimates of suffering and moral theories.

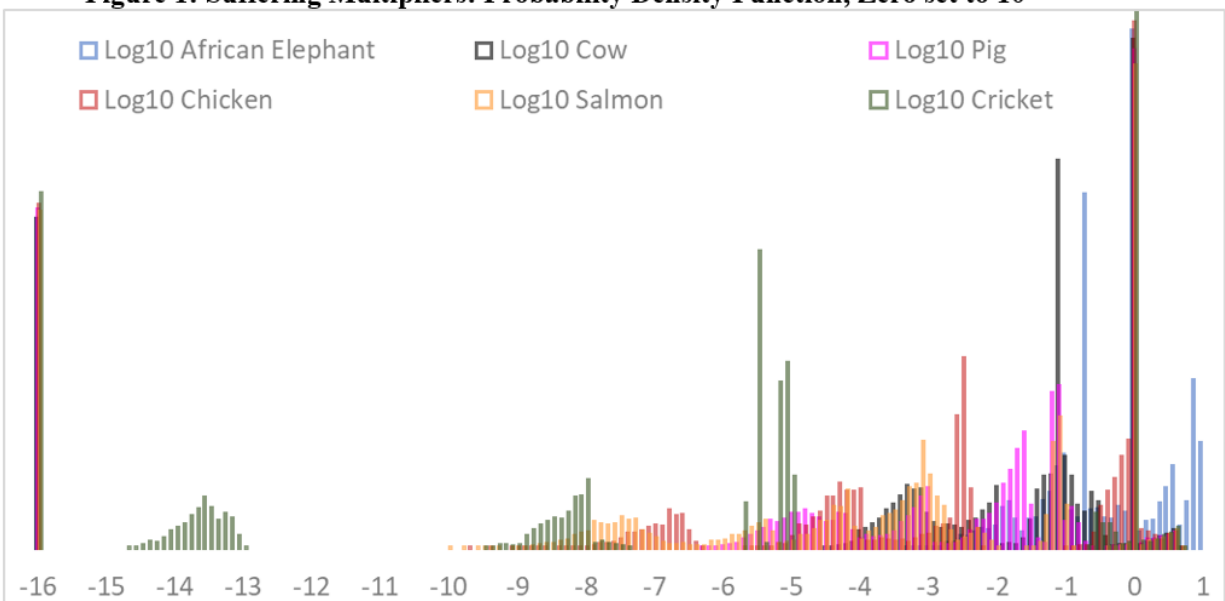
In each iteration, the procedure is as follows:

- 1) Choose a suffering theory, based on the credence. For example, the 'Unconscious Animal' theory is chosen in 3% of the iterations.
- 2) Choose a moral theory, based on the credence. For example, the 'Extreme Speciesism' theory is chosen in 10% of the iterations.
- 3) Choose the values of the theory's inputs, if any, from the specified range or probability distribution (e.g., brain and body mass).
- 4) Calculate the final moral value based on the randomly chosen theories and input values.
- 5) If the final value is zero, change it to a very small number of our choice (10^{-16}) for the purpose of graphing on a log scale and calculating the geometric mean.

We can graph the results of these thousands of iterations and see the spread of uncertainty in animal ethics. All graphs are on a logarithmic scale. We ran the simulation in Microsoft Excel (v16.37 for Mac) using the YASAI plug-in⁶ to run Monte Carlo Simulations. The spreadsheet is available at <http://www.jimdavies.org/science-of-better/>

Results

Figure 1: Suffering Multipliers: Probability Density Function, Zero set to 10^{-16}



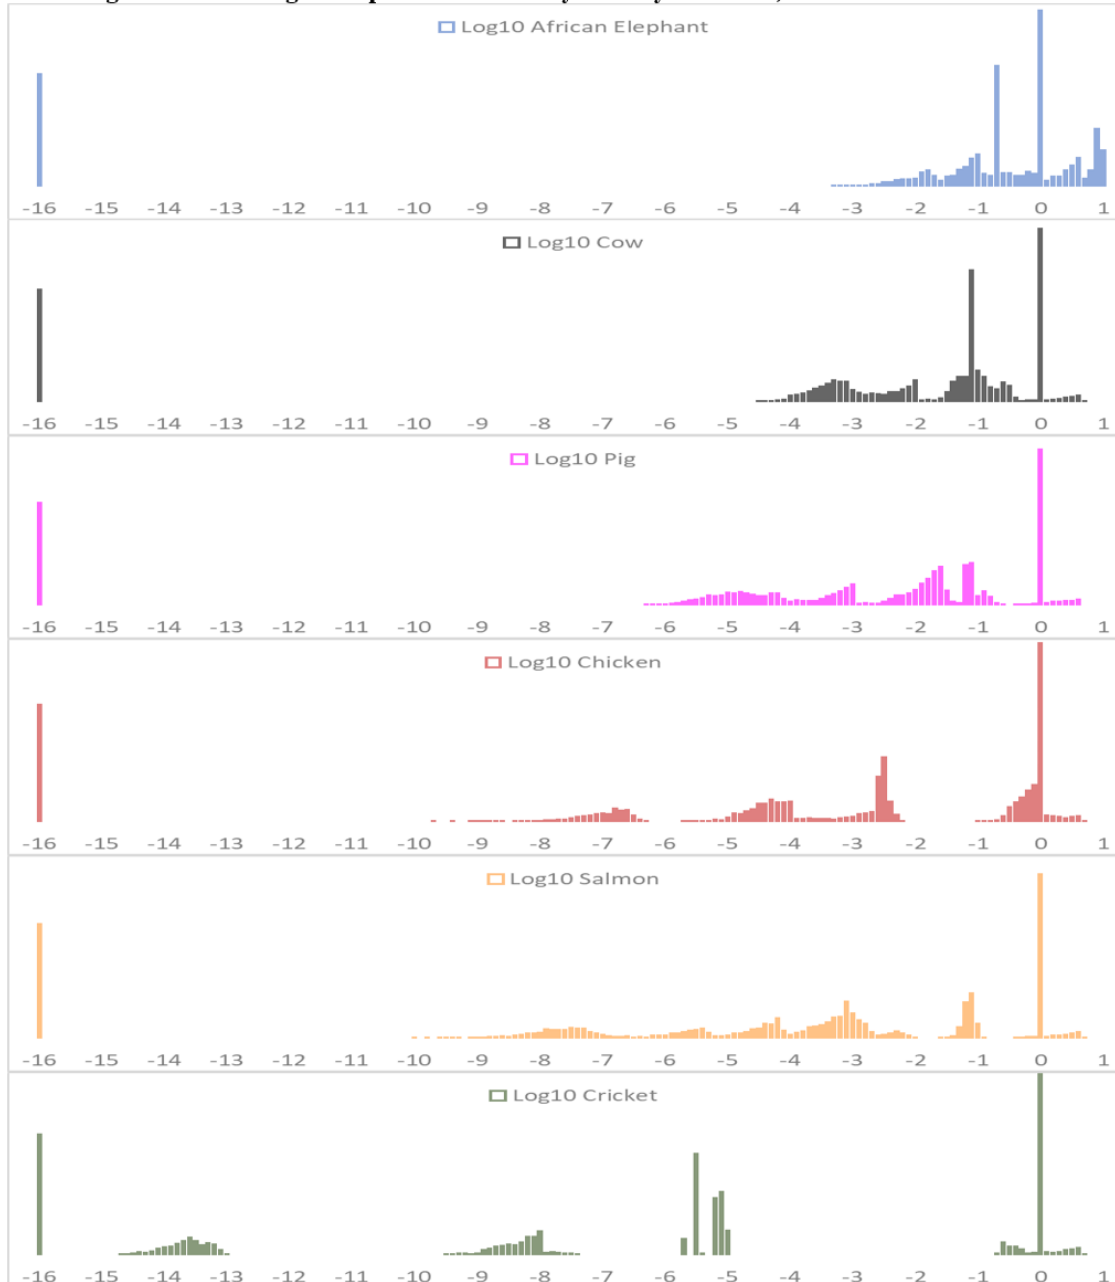
The height of the bars is the number of simulations that result in a certain value, which corresponds to the probability that the true value is that number. The horizontal axis is a log scale, for example -3 means 10^{-3} or 0.001. $10^0=1$ and $10^1=10$. We see a shift in probability distribution as animals go from more to less complex.

⁶ This plug-in is available free at <http://yasai.rutgers.edu/downloads.html>



Fig. 1 shows all animals on the same graph, and Fig. 2 shows each of the animals on their own graphs. We consider these figures, and the uncertainty they demonstrate, to be the central result, and recommend that any measures of central tendency only be used with extreme caution.

Figure 2: Suffering Multipliers: Probability Density Function, Zero set to 10^{-16}



The height of the bars is the number of simulations that result in a certain value, which corresponds to the probability that the true value is that number. The horizontal axis is a log scale, for example -3 means 10^{-3} or 0.001. $10^0=1$ and $10^1=10$. We see a shift in probability distribution as animals go from more to less complex.



Because the 90% confidence interval for all animals other than the elephant is basically 'zero to one', we present much smaller confidence intervals than are usually reported. This is similar to reporting the median value, as discussed below, while still conveying some of the uncertainty. However, as with reporting the median, it ignores potentially important differences in the magnitudes of the extreme values.

Table 1: 20% and 40% confidence intervals for weights of animal suffering. These numbers would be multiplied by some number representing the moral state or change to a human to get an estimate of the moral state or change for an animal.

Animal	20% Confidence Interval		40% Confidence Interval	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound
African Elephant	$2 \cdot 10^{-1}$	1	$7 \cdot 10^{-2}$	1
Chicken	$9 \cdot 10^{-5}$	$3 \cdot 10^{-3}$	$3 \cdot 10^{-5}$	$5 \cdot 10^{-1}$
Cow	$9 \cdot 10^{-3}$	$8 \cdot 10^{-2}$	$1 \cdot 10^{-3}$	$1 \cdot 10^{-1}$
Cricket	$1 \cdot 10^{-8}$	$6 \cdot 10^{-6}$	$2 \cdot 10^{-9}$	$9 \cdot 10^{-6}$
Pig	$9 \cdot 10^{-4}$	$2 \cdot 10^{-2}$	$7 \cdot 10^{-5}$	$7 \cdot 10^{-2}$
Salmon	$6 \cdot 10^{-5}$	$1 \cdot 10^{-3}$	$4 \cdot 10^{-6}$	$6 \cdot 10^{-2}$

Three candidates for a point estimate are the arithmetic mean, the geometric mean, and the median.

The arithmetic mean is greatly skewed by the Tinker Bell and Same Suffering theories, when combined with the Equal Consideration theories. Given that many results will output a moral value of one, the minimum value of the arithmetic mean will be about $1/3^{\text{rd}}$. More complex animals will have a slightly higher value. Therefore, we do not recommend using this value.

The geometric mean is often a more appropriate measure of central tendency when things are multiplied together, and our results are the suffering weight times the moral weight. The geometric mean also has the property that it is not unduly influenced by the high end of the distribution, in a situation where the results span many orders of magnitude. However, we must replace zeroes with very small values, and the choice of the small value matters. This semi-arbitrary choice could change the mean by several orders of magnitude.

The median does not have some of the problems associated with the arithmetic and geometric means, but that it throws away any differences in the low and high values. It essentially ignores the results of moral theories that produce high or low values. And there is also the possibility that, with different credence values, 50% of the simulation runs could have been an extreme value like zero or one. In this case, the median would ignore whatever was said by the minority theories.



Table 2: Measures of Central Tendency. These numbers represent moral weights for different animals, when taking into consideration the uncertainty as presented in the literature.

Animal	Arithmetic Mean	Geometric Mean	Median
African Elephant	1	5×10^{-3}	2×10^{-1}
Chicken	3×10^{-1}	6×10^{-5}	2×10^{-3}
Cow	3×10^{-1}	6×10^{-4}	6×10^{-2}
Cricket	3×10^{-1}	2×10^{-7}	3×10^{-6}
Pig	3×10^{-1}	1×10^{-4}	1×10^{-2}
Salmon	3×10^{-1}	2×10^{-5}	5×10^{-4}

Conclusion

Although we suspect many readers already knew that our understanding of animal suffering and ethics has a good deal of uncertainty, one thing we want to make clear in this paper is just how great the uncertainty actually is, if we take seriously the variety of viewpoints experts have on the subject. But what are we to do with this more precise analysis of the uncertainty?

First, we provide new point estimates and ranges that take into account the uncertainty. These can, with caveats, be used in utilitarian calculations when a number or range is needed. Although the uncertainty we present is large, it is still an improvement on the existing state of knowledge, and can be temporarily used in a ‘something is better than nothing’ sense until better estimates become available. Conducting surveys to get percentages of the endorsement of the suffering and moral theories is an important next step, and might change the results of these simulations significantly.

Second, if one is concerned about one’s own ethical behavior, or government or corporate policy, these uncertainties can be compared to other uncertainties, such as those involved with improving human global health or preserving the environment. For example, though there is disagreement on how much humans suffer and what can be done about it, there is much more agreement that (nearly) all humans can suffer greatly and equally, that reducing that suffering is a good thing, that (nearly) all humans should get the same moral treatment, and so on. When deciding whether to try to improve the state of animals or human beings, the fact that there is *greater* uncertainty regarding the good that can be done for animals might make human intervention more defensible.



Limitations of this work

This work depends on numerical calculation. Some of the numbers are very certain, such as human mass. Other numbers have much less certainty, such as the number of neurons in a cow brain. Worst are the numbers we simply guessed at: the credence levels of the suffering and moral theories, for instance. This is particularly problematic because the results are very sensitive to these credence levels. We want to emphasize that many of the estimates we use in this paper should be updated with better ones as more data become available.

An additional limitation is that expert views on theories of suffering and theories of value are likely to be correlated. Therefore, after survey data becomes available, the simulation structure should be changed to sample a random expert, and take that expert's credence on the theories of suffering and value, rather than aggregating the credences and choosing them independently.

Gardener Comments

Heather Browning:

This article contains something promising, but it needs a bit more work to flesh it out properly. The use of a Monte Carlo simulation to assess confidence intervals for the estimates could be useful, but I worry a bit about the current inputs to the model. For instance, no details are provided about the literature survey they conducted to determine how many scholars endorsed the different views – how many papers were surveyed, where were they found, how were the inferences about the authors' views performed? Ideally, a full list of the papers surveyed should be provided in the appendix. It would be much stronger if the ranges of views were taken from a survey of selected experts, rather than simply from a reading of the literature (as only a small portion may be publishing their views). This is not something the authors need to do necessarily for this paper, but an acknowledgement of the limitations of their current sampling method would be a good idea.

The introductory sections are quite uneven – some theories are mentioned in only a few sentences, while others have several paragraphs of argument for their plausibility. It would also be helpful for the authors to provide more discussion of their results, what they say about the methods, how they relate to the theories discussed earlier in the paper, and what some of the next steps should be (e.g. how would the authors suggest their numbers are used? What methods would they use to refine the current estimates?)

A few specific comments:

Page 3, paragraph 3 – this seems like a strange example – a cow is only more psychologically complex than a mouse on a small number of the metrics suggested; perhaps use something like a frog as a better contrast (**note**: this comment was addressed and frog was substituted for cow). Additionally, it would be useful to provide



some more detail on why these seem like plausible measures of psychological complexity (and, if possible, citation to where they are advocated in the literature)

Page 3, paragraph 6 – the inference that the relationship between human infant and human adult (in terms of the ‘pain multiplier’) is a good characterisation of that between human adult and animal is not well-supported here.

Page 4: “This relationship between intelligence and sentience is assumed by the Muted Intelligence theorists.” – there are several scholars who would explicitly argue against this, in favour of a hedonic capacity, unrelated to intelligence

Anonymous1:

The paper tries to formalize intuitions and justifications for the moral weights of certain animals. It is obvious that the paper is highly speculative, with the estimates used in the Monte Carlo simulation seemingly pulled out of thin air, so it does not provide much more value than simple intuition. But it is an interesting approach, and with more data it would be valuable to have quantified uncertainty in viewpoints.

Simon T. van Baal:

I think this paper is doing some important work, and it is highly interesting, but I have several suggestions. I will start with my most important comment; the rest are not essential. The first comment could make the simulations more useful, and might actually make them converge on a more precise estimate.

1) As I see it, the most important improvement the authors could make is to make the assumptions for the simulations more realistic by including varying conditional probabilities between views. To illustrate, it is unlikely that someone who is an Extreme Speciesist is just as likely to also hold the Tinker Bell view as someone who holds the Equal Consideration view. The solution, I think, is to either i) guess the conditional probabilities of holding one view for suffering, while holding another view for moral value, or preferably ii) collect some preliminary data on these conditional probabilities in the philosophy community through e.g. mailing lists, conferences, or social media. To me it seems the latter approach would enable the paper to make a larger step.

2) The authors assume a 1:1 relationship between suffering and the statistic the multiplier is based on (e.g., encephalisation quotient). This needn't be the case, however. Taking the case of the cow, there could - and probably should - be some (non-linear) transformation of the ratio in encephalisation quotient. I think merely stating the assumption would be sufficient here.

3) The first paragraph of the Theories of Value section is slightly confusing, because it appears the authors are not addressing how different 'events' might impact different



agents differently, but whether their suffering matters morally. To me this could be clearer, but I might have missed something.

4) Though I appreciate the detail in the figures, two figures based on bar plots depicting the same information are not the best way to depict this sort of data with many groups. I would suggest using separated raincloud plots instead or just to remove figure 1.

R. Sal Reyes:

This is an interesting approach to providing some measurable ways to estimate the perceived impact of (and how we “should” view) human actions upon other animals. And the framework seems broad enough in its considerations to easily include new or different factors in its simulations—either adding to, revising or replacing some of the current categories to help provide a deeper or clearer picture of the possible impact (or presence) of “suffering” in non-human animals.

One such addition or revision might involve, for example, the way that the framework scales the different possible levels of experienced suffering within a “Muted Consciousness” spectrum. As the paper notes: “How best to measure the relevant psychological complexity is also a matter of debate.” The paper’s current method for measuring that complexity—essentially equating a greater capacity for “intelligence” with a capacity for greater or more impactful suffering, and judging that intelligence according to measures of neural populations & mass in different animals—is a fairly straightforward approach, but it may not be as revealing a measure as it seems.

If we look at suffering as the purely sensory experience that is at the heart of it, and then view that sensory experience through the same lens that we view other sensory experiences (like sound or smell), it doesn’t seem that a more “complex” (aka, more intelligence-enabled) sensory experience should necessarily equal a more *intense* sensory experience. In other words, a more complex auditory cortex might enable a sound experience that has more detail and depth, but that complexity doesn’t likely enhance that brain’s ability to produce a very loud (aka, intense) sound. It seems more likely that a brain’s capacity to produce a wide range of differently intense sensory response experiences—like specific kinds of loud sounds or powerful smells—is present (and *necessary* for varied functionality) in any animal that possesses the capacity for that specific kind sensory experience (like a specific kind of pain-based experience that humans would describe as suffering), regardless of whether or not that brain has yet evolved more robust neural capacities that enable that specific sensory experience to become more complex & nuanced.

If the framework for these simulations adopted this “systematic” view in place of its current method for scaling “Muted Consciousness” (or included a separate category for it) then it might instead use the presence of specific neural systems and specific connectivity between these systems as ways to identify or scale different animals’ capacity for a high-impact or high-intensity sensory experience that results from or



defines or accompanies perceived suffering. Doing so would bring into these simulations an accounting for the *purpose* of suffering, because any new specific sensory capacity (& its supporting neural architecture) ought to have an adaptive functional purpose if it is to persist & become more complex.

This means that this sensory experience (suffering) must have a functional impact on behavior—helping to shape that behavior in some way that is useful in responding to the source of the suffering. Thus, if we identify the purpose of the suffering and are able to tie that perceived (or hypothesized) sensory experience to specific behavior that aids in fulfilling that purpose, we can use the observations of specific kinds of behaviors in different animals as kinds of markers indicating the likely presence of those specific neural sensory-experience systems. This would allow the framework to scale different levels of “Muted Consciousness” by determining what the animal’s likely experience is according to the specific behaviors it exhibits. (Instead of measures of its neural populations & mass.)

This kind of method might allow for a more nuanced middle ground between the "Same suffering" black & white view that says any feeling means all the same feeling, and the current "Muted Consciousness" experiential scale based on brute computing capacity. For example, using the behavioral method of scaling, a particular species of rodent may be observed displaying behavior that correlates specifically to basic pain responses, but that behavior does not appear to serve any of the purposes of perceived suffering (& its correlating behavior) identified in "higher level" animals like humans. This particular rodent might then be measured as having the capacity for feeling pain, but not experiencing a high level of "suffering."

To put it in the smallest nutshell I can manage: if there is a need for suffering, then there is likely the presence of suffering.

Dan James:

This is an attempt to quantify suffering in a few selected animals using a statistical analysis based on different moral conceptions. Though the authors freely admit to adopting a utilitarian approach, it seems to me that fundamental to their whole project is the moral status of animals, a question they themselves pose in the introduction. What the authors do not point out is that we have no good reason to assume this moral status can be understood solely within a utilitarian framework and indeed it is clearly open to interpretation in other ways, for example with deontological animal ethics - ‘The Case for Animal Rights’ Tom Regan 1983. (Though deontological animal ethics itself can have severe practical problems, even absurdities in terms of absolutist tendencies).

So for me this ‘meta’ level of deciding the moral status of animals is crucial, for if we concentrate or promote a biased interpretation, subsequent efforts may be futile. I appreciate of course that this might have led to a much lengthier and different kind of paper though.



The authors deftly avoid the paper becoming a dry accounting exercise by the inclusion of a broad and fascinating list of variables, some of which are claimed as unique to this paper, for example the inspired 'Tinker Bell Theory' (which is so good it deserves to be true!)

Overall I feel there is a danger in using statistics that are based on some fairly generalised assumptions about how to classify animal suffering, and to draw the conclusion that such surveys are 'better than nothing'. The danger here being that poor welfare standards or even inflicted harm can conceivably be legitimised or gain credence by a purported metric of suffering.

Despite these misgivings I felt it was a fascinating paper with complex statistical analysis presented in a very understandable way and consequently have no hesitation in recommending publication.

Michael M. Kazanjian:

I found this to be a brief but strong, stats based evaluation of animal protection. It is brief but sufficiently well documented to inform the reader of the details involved in a utilitarian approach. It adds well to the literature.

William Collen:

It does not appear that this article meets the first criterion for inclusion in Seeds of Science. The data collected in the article might be of use in the domain of ethical philosophy, but not, in any way that I see, in the realm of scientific inquiry in general. The author of the paper proposes a novel method of exploring the topic, but the method uses previously-known and accepted methods of scientific research, and the scope of the paper's method is too narrow to count as an advancement of scientific thought in any way other than the narrow application presented by the author.

Sam Harsimony:

It would be nice to see an additional column in table 2 comparing these results to an outside source such as a survey of human experts.

In the future, the author may want to consider using variance normalization to weight different moral theories ("[Geometric reasons for normalising variance to aggregate preferences](#)").

Mark:

This is interesting work that could be rather impactful, however, I think it's missing some important parts which will lead to it being disregarded by some readers and communities.



The main factor I think reduces the potential impact of the work is that although it seems aware of the related existing work it seems to cite it sparsely, even when openly sliding to ideas from that broader literature. As a consequence the current writing seems poorly motivated. With a bit of work making these references more consistent and expanding them to have full coverage, this work could be more deservedly regarded.

I also think the clarity and quality of the figures could be improved, making them more clear but also better communicate the modeling assumptions.

Roger's Bacon:

This article represents a useful (and to my knowledge) novel approach to an age-old question and for that reason I believe it deserves to be published, however, I believe there are serious issues that need to be addressed in revisions before publication. The primary issue is that the numbers used for input into the MCMC simulation are poorly justified (if they are justified at all). I think that in some cases a more thorough review of the literature will yield better reasons/evidence for the chosen numbers, but in the end there will still be a great deal of subjectivity - that's ok if the article is couched as a proof-of-concept analysis but I find fault with the "something is better than nothing" conclusion as this will decidedly not be true if the input to the simulation is poorly chosen (which I think many people will say that it is). I wonder if this methodology might work better for other philosophical issues which can provide higher-confidence estimates of the underlying quantitative measures.

Some more specific feedback:

What about modeling these theories as all equally probably - i.e. maximal uncertainty?

- While I appreciate the challenge of what the authors are trying to do, the whole meditation and newborn pain section feels completely apples to oranges to me and basically useless as far as providing info for simulation inputs. I think there needs to be a lot more discussion of the theoretical and empirical issues here (or just none at all and say that frankly we chose numbers that seemed reasonable to us). Wouldn't behavioral/intelligence measures be more useful here - so just the inverse of the inputs to the muted animal theory - i.e. if dogs are 20% smart as humans then we would say they suffer 5x compared to humans under Tinker Bell theory?

References

Akhtar, S (2011) Animal pain and welfare: Can pain sometimes be worse for them than for us. *The Oxford handbook of animal ethics*, 495-518.



Azevedo FA, Carvalho LR, Grinberg LT, Farfel JM, Ferretti RE, Leite RE, ... & Herculano-Houzel S (2009) Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Neurology*, 513(5), 532-541.

Ballarin C, Povinelli M, Granato A, Panin M, Corain L, Peruffo A, Cozzi B (2016) The brain of the domestic *Bos taurus*: weight, encephalization and cerebellar quotients, and comparison with other domestic and wild Cetartiodactyla. *PloS one*, 11(4).

Burger CV, Wilmot RL, Wangaard DB (1985) Comparison of spawning areas and times for two runs of chinook salmon (*Oncorhynchus tshawytscha*) in the Kenai River, Alaska. *Canadian Journal of Fisheries and Aquatic Sciences*, 42(4), 693-700.

Carruthers P (2000) *Phenomenal Consciousness*. Cambridge: Cambridge University Press.

Carruthers P (2011) Animal mentality: Its character, extent, and moral significance. In T.L. Beauchamp and R.G. Frey (Eds). *The Oxford Handbook of Animal Ethics*. Oxford, UK: Oxford University Press.

Clarke TC, Barnes PM, Black LI, Stussman, B. J., & Nahin, R. L. (2018) *Use of yoga, meditation, and chiropractors among US adults aged 18 and over*. US Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Health Statistics.

Copp D (2011) Animals, fundamental moral standing, and speciesism. In T.L. Beauchamp and RG Frey (Eds). *The Oxford Handbook of Animal Ethics*. Oxford, UK: Oxford University Press.

Crile G and Quiring DP (1940) A record of the body weight and certain organ and gland weights of 3690 animals. *The Ohio Journal of Science*. XL(5), 219-259.

DeGrazia D (2008) Moral status as a matter of degree?. *The Southern Journal of Philosophy*, 46(2), 181-198.

Dennett DC (1995) Animal consciousness: What matters and why. *Social Research*, 691-710.

de Waal FB and Ferrari PF (2010) Towards a bottom-up perspective on animal and human cognition. *Trends in cognitive sciences*, 14(5), 201-207.

Jelsing J, Nielsen R, Olsen AK, Grand N, Hemmingsen R and Pakkenberg B (2006) The postnatal development of neocortical neurons and glial cells in the Göttingen minipig and the domestic pig brain. *Journal of experimental biology*, 209(8), 1454-1462.

Harrison P (1991) Do animals feel pain?. *Philosophy*, 66(255), 25-40.



Herculano-Houzel S (2009) The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in human neuroscience*, 3, 31.

Herculano-Houzel S, Catania K, Manger PR and Kaas JH (2015) Mammalian brains are made of these: a dataset of the numbers and densities of neuronal and nonneuronal cells in the brain of glires, primates, scandentia, eulipotyphlans, afrotherians and artiodactyls, and their relationship with body mass. *Brain, Behavior and Evolution*, 86(3-4), 145-163.

Huebner B (2011) Minimal minds. In T.L. Beauchamp and RG Frey (Eds). *The Oxford Handbook of Animal Ethics*. Oxford, UK: Oxford University Press.

Goksan S, Hartley C, Emery F, Cockrill N, Poorun R, Moultrie F., ... and Clare S (2015) fMRI reveals neural activity overlap between adult and infant pain. *Elife*, 4, e06356. <https://elifesciences.org/articles/06356>

Kagan S (2019) *How to Count Animals, More or Less*. Oxford, UK: Oxford University Press.

Kazu RS, Maldonado J, Mota B, Manger PR and Herculano-Houzel S (2014) Cellular scaling rules for the brain of Artiodactyla include a highly folded cortex with few neurons. *Frontiers in Neuroanatomy*, 8, 128. [doi:10.3389/fnana.2015.00039](https://doi.org/10.3389/fnana.2015.00039)

Key B (2015) Fish do not feel pain and its implications for understanding phenomenal consciousness. *Biology & philosophy*, 30(2), 149-165.

Koller D and Goldman RD (2012) Distraction techniques for children undergoing procedures: a critical review of pediatric research. *Journal of pediatric nursing*, 27(6), 652-681.

Lisney TJ and Collin SP (2006) Brain morphology in large pelagic fishes: a comparison between sharks and teleosts. *Journal of Fish Biology*, 68(2), 532-554.

Lutz A, McFarlin DR, Perlman DM, Salomons TV and Davidson RJ (2013) Altered anterior insula activation during anticipation and experience of painful stimuli in expert meditators. *Neuroimage*, 64, 538-546.

MacAskill W (2014) 'Normative Uncertainty' (DPhil thesis, Oxford University, 2014).

Malda J, de Grauw JC, Benders KE, Kik MJ, van de Lest CH, Creemers LB, ... and van Weeren, P. R. (2013) Of mice, men and elephants: the relation between articular cartilage thickness and body mass. *PloS one*, 8(2).



Matsumoto Y (2018) Structure of the neuropils in the brain of the two-spotted cricket. https://invbrain.neuroinf.jp/modules/htmldocs/IVBPF/Cricket/Cricket_brain.html Accessed June 16, 2020

Morris S, Humphreys D and Reynolds D (2006) Myth, marula, and elephant: an assessment of voluntary ethanol intoxication of the African elephant (*Loxodonta africana*) following feeding on the fruit of the marula tree (*Sclerocarya birrea*). *Physiological and Biochemical Zoology*, 79(2), 363-369.

Minervini S, Accogli G, Pirone A, Graic JM, Cozzi B and Desantis S (2016) Brain mass and encephalization quotients in the domestic industrial pig (*Sus scrofa*). *PloS one*, 11(6).

Mink JW, Blumenshine RJ and Adams DB (1981) Ratio of central nervous system to body metabolism in vertebrates: its constancy and functional basis. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 241(3), R203-R212.

Olkowicz S, Kocourek M, Lučan RK, Porteš M, Fitch WT, Herculano-Houzel S and Němec P (2016) Birds have primate-like numbers of neurons in the forebrain. *Proceedings of the National Academy of Sciences*, 113(26), 7255-7260.

Panigrahy KK, Behera K, Mandal AK, Sethy K, Panda S, Mishra SP, ... and Sahu SK (2017) Gross morphology and encephalization quotient of brain in male and female Vanaraja chickens at different ages. *Exploratory Animal and Medical Research*, 7(1), 28-32.

Roth G (2013) Comparing vertebrate brains. In *The Long Evolution of Brains and Minds* (pp. 223-242). Springer, Dordrecht.

Roth G, and Dicke U (2005) Evolution of the brain and intelligence. *Trends in cognitive sciences*, 9(5), 250-257.

Sandhoefner E (2017, June 11) Animal Suffering Calculator. Retrieved from Animal Suffering Calculator website: http://sandhoefner.github.io/animal_suffering_calculator Accessed June 16, 2020.

Scherer L, Tomasik B, Rueda O and Pfister S (2018) Framework for integrating animal welfare into life cycle sustainability assessment. *The international journal of life cycle assessment*, 23(7), 1476-1490.

Schubert H, Wood S, Reyher K and Mills H (2019) Data-driven approach to using individual cattle weights to estimate mean adult dairy cattle weight. *Veterinary Record*. doi:10.1136/vetrec-2019-105471



Sergo P (2008) Wasps reveal clues about the evolution of intelligence. *Scientific American Mind*, June 1.

<https://www.scientificamerican.com/article/big-brains-dominate/> Accessed June 16, 2020

Shriver A (2009) Knocking out pain in livestock: Can technology succeed where morality has stalled?. *Neuroethics*, 2(3), 115-124.

Thomson JJ (1990) *The realm of rights*. Cambridge, MA: Harvard University Press.

Tomasik B (2018, July 14) How Much Direct Suffering Is Caused by Various Animal Foods? Essays on Reducing Suffering website:

<https://reducing-suffering.org/how-much-direct-suffering-is-caused-by-various-animal-foods/> Accessed June 16, 2020

de Waal FB and Ferrari PF (2010) Towards a bottom-up perspective on animal and human cognition. *Trends in cognitive sciences*, 14(5), 201-207.

Wiper ML, Britton S and Higgs DM (2014) Early experience and reproductive morph both affect brain morphology in adult male Chinook salmon (*Oncorhynchus tshawytscha*). *Canadian Journal of Fisheries and Aquatic Sciences*, 71(9), 1430-1436.

Zeidan F, Martucci KT, Kraft RA, Gordon NS, McHaffie JG and Coghill RC (2011) Brain mechanisms supporting the modulation of pain by mindfulness meditation. *Journal of Neuroscience*, 31(14), 5540-5548.

Supplementary Material

The simulations are done on a Microsoft Excel spreadsheet. One page has estimates for humans, and there is a separate page for each of the six other animals we compare. The YASAI plugin, which is free to download from <http://yasai.rutgers.edu/downloads.html>, is required to run the simulations. Download the plugin to the desktop, open the spreadsheet, then activate the plugin with Tools -> Excel Add-Ins... Navigate to the file on your desktop and open it.

Appendix: Sources of Biological Estimates

For all of the estimates required for the simulation, we tried to find high and low estimates for the simulation to sample from. In this section we will describe the sources. All mass is presented in grams.

Human

The high and low estimates for human body mass come from 95% confidence intervals given by WolframAlpha.com: 49,000g to 128,000g (49kg to 128kg, or 108 pounds to 282 pounds). Human brain mass ranges from 1209.77g to 1808.05g (Azevedo et al.



2009). This means the human brain-body ratio ranges between 0.011 and 0.026. The human Encephalization Quotient ranges from 7.4 to 7.8 (Roth and Dicke 2005). The number of neurons in the human brain is about 86.06 billion, plus or minus 8.12 billion, or 78,820,000,000 to 95,400,000,000 neurons (Azevedo et al. 2009). The human cortex contains between 11,500,000,000 (Roth and Dicke 2005) and 18,510,000,000 (Azevedo et al. 2009) neurons.

Cricket

There are no estimates we could find for insect Encephalization Quotients. For body and brain mass, we assume it is the same as a grasshopper: a body weight of 0.46g and a brain mass of 0.0023g (Crile and Quiring 1940, 221). This means a brain-body ratio of 0.2. For neuron counts we use Burne et al. (2011)'s fruit fly and ant models to estimate that there are 250,000 neurons in a cricket. A cricket does not have a neocortex, so the lower estimate is 0. There are about 100,000 neurons in the part of the cricket brain called the "mushroom body." Some have suggested that the mushroom bodies of crickets are functionally equivalent to mammalian cerebral cortexes (Sergo 2008), so we use 100,000 neurons as the upper estimate of cortical neurons in a cricket (Matsumoto 2018).

Salmon

We eat many kinds of fish, and they range greatly in size. We use chinook salmon as a representative fish, as it is commonly eaten, and because there is a lot of data available about them. Salmon are ray-finned fish (as opposed to cartilaginous fish, like sharks and rays). Salmon brains are estimated to be between 1 and 1.2g. These numbers are from eyeballing Fig.4A of Wiper, Britton and Higgs (2014). The estimated weight of a chinook salmon is 18kg (U.S. Fish & Wildlife Service 2020). If we assume the range is from 50% of that to 150% of that average, we get a lower estimate of 9000g and a higher estimate of 36000g. This is close to estimates given elsewhere (8-32kg in Burger, Wilmot and Wangaard 1985). Encephalization Quotients for ray-finned fish are estimated to be about 0.5 (This number comes from eyeballing a graph in Lisney and Collin 2006; we assume a variance of 0.1). To estimate the number of neurons in a chinook salmon, we use the mass of the brain, and then estimate the number of neurons based on whether it scales like a rodent, primate, or human (Herculano-Houzel 2009). Using the low brain weight of 1g, and the smallest (rodent at 8 million per gram) scaling, we get a lower estimate of 8000000 neurons. Using the high brain weight of 1.2g, and the largest scaling (human at 57 million per gram), we get our higher estimate of 68400000 neurons. Fish do not have a neocortex, so our low estimate for neocortical neurons is 0. They do, however, have a forebrain, which is approximately 1/3 of their entire brain, giving us a higher estimate of 22800000.

Chicken

Chicken mass ranges from about 43.7g (Olkowicz et al. 2016, 224) to 1800g (Mink 1981, 28). Chicken brains have a masses ranging from 0.1g to 3.92g (Olkowicz et al. 2016, 224). The domestic chicken Encephalization Quotients are approximated from a Vanaraja chicken model, ranging from 1.8 to 6.14 (Panigrahy et al. 2017). For the



number of neurons in the domestic chicken brain, we use a red junglefowl model, which ranges from 190,809,468 to 271,959,440 neurons. Birds do not have a neocortex, so our lower estimate for number of cortical neurons is 0, but as some believe that the pallium serves a similar function, and possibly evolved from the same reptilian telencephalic structure (de Waal and Ferrari 2010), we use the pallium neuron count of 61,000,000 neurons for our higher estimate of cortical neuron count (Scherer, Tomasik, Rueda and Pfister 2018).

Pig

Pig (*Sus scrofa domesticus*) ranges in body mass from 51550g (Mink 1981, 28) to 113200g (Crile and Quiring 1940). Pig Encephalization Quotient ranges from 0.38 to 0.58 (Minervini et al. 2016, 12). We only have one estimate of pig neuron count, at 307,000,000 (Herculano-Houzel 2016, Fig. 4.15). For cortical neuron counts, our lower estimate is 39,000,000 (13% of 303 million, Herculano-Houzel 2016) and the upper bound is 432,000,000 (Jelsing, Nielsen, Olsen, Grand, Hemmingsen and Pakkenberg, 2006, 1458).

Cow

Cow mass ranges from 413,000g (Crile and Quiring 1940) to 617,000g (Schubert, Wood, Reyher and Mills 2019). Cow brain mass ranges from 408g (Crile and Quiring 1940, 251) to 420g (Mink 1981, 28). The Encephalization Quotient of cows ranges from 0.58 to 0.59 (Ballarin Et al., 2016). We do not know how many neurons are in a head of female cattle, so we make estimates by scaling according to rodents and humans (as we did for salmon) to get lower and upper bounds of 3,264,000,000 to 23,940,000,000 neurons (Herculano-Houzel 2009). For cortical neurons, we get a low estimate by assuming similarly-weighted springbok models (*Antidorcas marsupialis*) at 375,490,000 (Kazu, Malonado, Mota, Manger, and Herculano-houzel 2014, 5) or a giraffe models (*Giraffa Camelopardalis*) at 1,730,000,000 (Herculano-Houzel, Catania, Manger and Kaas 2015).

African Elephant

The mass of an African Elephant ranges from 2,500,000g to 6,000,000g (Morris, Humphreys and Reynolds 2006), and their brain mass ranges from 4200g to 6000g (Malda et al. 2013). The Encephalization Quotient is around 1.3 (Roth 2013). They have an estimated 21 to 354 billion neurons, using the rodent and human scaling of Herculano-Houzel (2009) also used for salmon above. There are an estimated 11,000 million neurons in the elephant cortex (Roth 2013).