

Risk Aversion in Wild Animal Welfare

William McAuliffe — Senior Research Manager and Editor



Last updated: November 13, 2023

Table of Contents

- Executive Summary.....2**
- Wild Animal Welfare Has High Expected Value.....3**
- Complex Cluelessness in Wild Animal Welfare.....5**
 - Measuring Aggregate Welfare..... 6
 - Predicting Ecosystem Change..... 8
 - Abundance Versus Sentience.....9
- Risk Aversion Helps Cope with Uncertainty.....10**
- Three Types of Risk Aversion.....12**
 - Outcome Risk Aversion..... 13
 - Difference-Making Risk Aversion.....15
 - Ambiguity Aversion..... 15
- The Value– and Rarity– of Robustness.....16**
 - Convergent Evidence..... 16
 - Divergent Recommendations.....18
 - Probability Versus Value..... 18
 - Benefits of Research.....18
- Risk Aversion Affects Animal Welfare Cause Prioritization.....19**
- Conclusion..... 21**
- Acknowledgments..... 23**
- References..... 23**

Executive Summary

1. Wild animals outnumber humans and captive animals by orders of magnitude. Hence, scalable interventions to improve the welfare of wild animals could have greater expected value than interventions on behalf of other groups.
2. Yet, wild animals receive only a small share of resources earmarked for animal welfare causes. This may be because animal advocates are uncomfortable with relying on expected value maximization alone in a field beset by "complex cluelessness": There are compelling reasons for *and* against wild animal interventions, and none are clearly decisive.
 - a. Reducing populations of fast life history strategists would likely reduce suffering. However, there is also reason to suspect fast life history strategists have enough rewarding experiences to increase aggregate welfare.
 - b. Eliminating fundamental sources of suffering in natural habitats would reduce suffering. However, it could also differentially benefit species that many people believe have systematically worse lives.
 - c. Prioritizing the most abundant groups of wild animals could generate the largest increases in aggregate welfare. However, the most abundant wild animals have relatively low and vague probabilities of sentience.
3. Regardless of risk attitudes, inaction on wild animal welfare is difficult to justify.
 - a. There are no larger-scale areas of animal welfare where less risky interventions might be found.
 - b. Even if the aggregate welfare of wild animals is net-positive, it is nevertheless almost certainly suboptimal.
4. By accounting for considerations that decision-makers believe are relevant, incorporating risk aversion into expected value calculations may increase willingness to commit resources to wild animal welfare. Different types of risk aversion account for different types of uncertainty.
 - a. *Outcome* risk aversion gives special consideration to avoiding worst-case scenarios.
 - b. *Difference-making* risk aversion gives special consideration to ensuring that actions are improvements on the status quo.
 - c. *Ambiguity* aversion gives special consideration to reducing ignorance and choosing actions that have predictable outcomes.
5. Different types of risk often disagree in their recommendations. A corollary is that robustness across different types of risk aversion increases choiceworthiness.



- a. Interventions that reduce suffering without altering the number or composition of wild animals have a greater probability of robustness to different types of risk aversion.
 - b. Outcome risk aversion favors abundant groups of wild animals, while difference-making risk aversion favors wild animals who have a high probability of sentience.
 - c. Ambiguity aversion is favorable towards research on wild animal welfare, whereas outcome and difference-making risk aversion only favor it under certain conditions.
6. Risk aversion does not robustly favor farmed over wild animals or vice versa.
 - a. Outcome risk aversion prioritizes wild animals due to their abundance.
 - b. Difference-making risk aversion favors farmed animals. However, it also favors some diversification across types of animals.
 - c. Ambiguity aversion favors helping farmed animals over wild animals, and basic research to help both groups.
 7. Although complex cluelessness affects many domains, wild animal welfare may be a particularly high-stakes example of it. Alternatively, moral uncertainty about the permissibility of interfering with nature may explain a reluctance to act on uncertain evidence.

Wild Animal Welfare Has High Expected Value

For those who want to impartially maximize welfare, helping wild animals looks like a promising bet. This is primarily due to their enormous abundance. To give a sense of scale, there are currently around [8.1 billion humans](#). [Bar-On et al., 2018](#) estimates that there are around 100 quintillion wild marine arthropods (e.g., krill, horseshoe crabs, sea spiders) alone (see Table S1). That means there are over twelve billion marine arthropods for every person.

Imagine you had a chance of helping marine arthropods or any one group of farmed animals. Which option has higher expected value? Expected value is the sum of the values of possible outcomes, weighted by their probabilities. Here, it is the probability that you will help a given group of animals multiplied by the value that results if you do help, plus the probability that you will not help multiplied by the value that results if you do not. The probabilities are determined by factors such as how likely it is that the animals in question are sentient, and how likely it is that the intervention to help them will work. The value is determined by factors such as the numbers of individuals affected and how much the intervention increases the welfare of each affected individual.



Suppose that the probability of successfully helping marine arthropods is just .0001, and the proportion of all marine arthropods helped is also .0001. The welfare of each successfully helped arthropod would increase by 1 unit. Nothing welfare-relevant happens if the intervention fails, so you can ignore that possibility in the calculation. The expected value is $(10^{20} \text{ marine arthropods} \times 10^{-4} \text{ of marine arthropods helped} \times 1 \text{ unit of welfare} \times 10^{-4} \text{ probability of success}) = 10^{12}$. For farmed animals, you get to choose which group of animals to help; you choose the most abundant taxon, shrimp, with 230 billion individuals at any given time ([Waldhorn & Autric, 2023](#)). The probability of helping is 1, a certain success. The proportion of farmed shrimp helped is also 1— each and every individual will be better off. As before, the improvement in welfare is 1 unit per individual. The expected value of helping farmed shrimp is $(2.3 \times 10^{11} \text{ farmed shrimp} \times 1 \text{ unit of welfare} \times 1 \text{ probability of success}) = 2.3 \times 10^{11}$. The sheer number of marine arthropods means that helping them is still a priority even if their probability of sentience is low or if the opportunity to help them is based on thin evidence. The welfare improvement per farmed shrimp would need to be higher for them to make up for their lesser abundance.

Similar examples would yield the same conclusion for other wild taxa, such as fishes and terrestrial arthropods. Yet, funding intended specifically for improving the welfare of wild animals is very low ([Animal Charity Evaluators, 2023](#)). Why are expected value considerations insufficient for prioritizing wild animal welfare? The present discussion is based on the hypothesis that stakeholders are risk averse. One piece of supporting evidence is that even organizations at the forefront of the field are not yet directly helping wild animals at a large scale or recommending that others do so. When doing public outreach, [Animal Ethics](#) cautions that "it may be useful in some very specific cases to mention other possible ways that could help larger numbers of animals but that are much more controversial. As a general rule, however, it might be better not to address these methods." [Wild Animal Initiative](#) emphasizes the need to develop *responsible* interventions when explaining their vision for a scientific field dedicated for wild animal welfare.

For risk-averse actors, high expected value alone is not a decisive reason to allocate a large share of resources to a field plagued with uncertainty. Stakeholders might nevertheless rely on expected value maximization to guide the *exploration* of cause areas because it is a useful framework for identifying high-value interventions. When the time comes to *implement* interventions, however, stakeholders may wish for principled alternatives that are tailored to their risk attitudes. Without the awareness of which interventions risk-aversion would recommend, the result may be inaction rather than compliance with risk neutrality or risk aversion.



To encourage stakeholders to overcome inaction, the present report illustrates how an explicit incorporation of risk aversion into decision-making would affect wild animal welfare strategy. Because resources are often earmarked for animal welfare, there is also an exploration of how risk aversion affects resource allocation between wild and farmed animal welfare. To frame the discussion, the report follows the lead of [Clatterbuck \(2023\)](#) and [Duffy \(2023c\)](#), who demonstrated how three types of risk aversion affect resource allocation decisions between interventions that benefit humans and those that benefit farmed animals.

Complex Cluelessness in Wild Animal Welfare

The advocacy organization [Animal Ethics](#) notes that humans already help wild animals on a small scale, such as saving them from natural disasters, providing them with shelter, and vaccinating them from disease. There is some uncertainty about whether these "small-scale" interventions do more harm than good. After being released, some aided animals may suffer an even more grisly death, live a long life characterized by unrelenting fear, or inflict more pain on prey or competitors than they otherwise would have experienced. Yet, the dangers that humans are averting in these cases are already so bad that the scope is wider for improving welfare than reducing it. Also, the scale of helping is typically too small to cause large negative side effects.

Small-scale interventions are generally not cost-effective in their own right. Enacting them might pave the way for large-scale interventions, by increasing the number of people dedicated to working on wild animal welfare, piloting new interventions that can be refined, and so on. At some point, though, wild animal interventions will need to affect a much larger number of individuals to be cost-competitive with farmed animal welfare. Many ideas for "large-scale" interventions involve systematic changes to entire species or the ecosystems they live in. Although they may have high expected value, they also introduce *complex cluelessness*. Complex cluelessness occurs when:

For some pair of actions of interest A1, A2,
 (CC1) We have some reasons to think that the unforeseeable consequences of A1 would systematically tend to be substantially better than those of A2;
 (CC2) We have some reasons to think that the unforeseeable consequences of A2 would systematically tend to be substantially better than those of A1;
 (CC3) It is unclear how to weigh up these reasons against one another.
 ([Greaves, 2016](#), p. 323)



Complex cluelessness can be contrasted with "simple" cluelessness, where there is no reason for believing that the indirect consequences of one action would be systematically better or worse than its alternative (ibid, p. 322). Under simple cluelessness, one can make expected value calculations on the basis of the direct¹ consequences alone even though the indirect consequences will likely have a greater collective impact on aggregate welfare. Indirect consequences are ignorable because there is no rational basis for predicting which alternative will have superior indirect consequences. When there is complex cluelessness, "one has more specific reasons for suspecting particular, systematic correlations between acts and 'indirect' effects, but too many such reasons: non-isomorphic reasons that point in different directions" (ibid, p. 336). These opposing considerations of unclear decision-making weight must somehow be factored into expected value calculations (ibid, p. 327), but it is "radically unclear" how they "should in the end be weighed against one another" (ibid, p. 325).

Most small-scale interventions do not face complex cluelessness to any significant degree. There are predictable positive and negative indirect effects, but the former clearly outweigh the latter. The mere fact that there *could* be unpredictable negative consequences would not provide a reason for hesitation. In contrast, large-scale interventions pose at least three issues that are beset by complex cluelessness.

Measuring Aggregate Welfare

Consider the widespread claim that wild animals experience more suffering than enjoyment in aggregate. The reason is that most individuals are members of species with "fast" life history strategies—roughly, they have very high birth rates to maximize the probability that at least a couple successfully reproduce. Over the long run, the vast majority of offspring will die prior to reproduction because the population cannot grow indefinitely without hitting some resource constraint. Allegedly, for those who die during the juvenile stage, there is little time to accumulate enough enjoyable experiences to compensate for a painful death ([Faria, 2023](#); [Ng, 1995](#); [Tomasik, 2015](#)).

The conclusion that the average experience of a wild animal is negative is based on a general principle of ecology rather than actual observations of suffering and pleasure.

¹ [Greaves \(2016\)](#) vacillates between contrasting "foreseeable" and "unforeseeable" effects and "direct" and "indirect" effects, respectively (e.g., p. 324, p. 336). The former can be confusing because the "unforeseeable" effects must be at least partly foreseeable for complex cluelessness to occur. The present report therefore refers to the most proximate, intended outcomes as "direct" effects, and all secondary outcomes, intended or unintended and foreseen or unforeseen, as "indirect."

Nevertheless, it has *some* force. But how much? A particularly difficult group to make generalizations about are terrestrial insect herbivores, which include aphids, most beetles, and most butterflies and moths. They "are often grouped near the 'r' or 'fast' end of life history classifications," ([Cuddington, 2019](#)) but this does not necessarily mean that they die young. The juvenile stage is often the longest life stage, sometimes taking years. In theory, many of them have ample opportunity for rewarding experiences. But, again, general observations raise more questions than they answer. How enjoyable is the daily life of a terrestrial insect herbivore? [Cuddington \(2019\)](#) observes that "juvenile insect herbivores seem to be seldom short of food." But how rewarding is ready access to food, especially when "food quality may vary"? And how common is it to live through most of the juvenile stage? A straightforward answer is elusive because there is "wide variation in survival trajectories regardless of whether the same species and population was sampled within the same year in different locations or between years."

Overall, high juvenile mortality rates are certainly *relevant* to whether wild animals are happy or unhappy in aggregate, but are not *decisive*.² There are also compelling reasons for thinking that individuals who die during the juvenile stage nevertheless have many rewarding experiences. Complex cluelessness complicates the assessment of interventions that reduce the number of fast life history strategists. To wit, some authors favor the proliferation of large herbivores such as elephants. The explanation is:

Elephants, unlike most wild animals, have large lifespans, very few offspring, high parental investment, and, as a consequence, high survival rates. They also consume a huge quantity of biomass, thus preventing both smaller animals with larger progenies from coming into existence and large trophic chains from emerging within an ecosystem. This means that ecosystems where elephants are present are likely to contain lower levels of suffering. ([Faria, 2023](#), p. 172).

From the perspective of reducing suffering, there is a compelling reason to preserve large herbivores. However, from the perspective of increasing the number of happy experiences, there is a reason to allow large herbivore populations to decline. The evidence backing one option does not systematically outweigh the evidence backing the other by a comfortable margin.

² You could argue that although high juvenile mortality rates are not decisive by themselves, the overall preponderance of evidence makes it clear that the aggregate welfare of wild animals is negative. However, [Browning and Veit \(2023\)](#) show that other common arguments for why suffering outweighs enjoyment also have limited evidential value.



Predicting Ecosystem Change

While some writers want to reduce suffering by reducing wild animal populations, others aim to reshape ecosystems so that they do not cause so much suffering. Two sources of suffering that are inherent to most natural habitats are predation and competition for resources. Many authors have discussed the possibility of using gene editing to eradicate both. Predator species would receive the traits required for herbivory. Altering predators only would cause prey populations to increase, which would exacerbate suffering due to competition. But gene editing (or some other technology, like birth control) could suppress reproduction to a level where all wild animals have ample resources ([McMahan, 2010](#)).

By addressing the root causes of suffering in natural habitats, altering the characteristics of wild species could have enormous welfare gains. Unsurprisingly, though, there are concerns about the indirect effects of such radical changes. [Delon and Purves \(2018\)](#) point out that species coevolve suites of traits that jointly function to promote inclusive fitness. Using gene editing to alter some traits without also altering coevolved traits may increase the likelihood of extinction. If the extinct species had fulfilled certain roles in an ecosystem (e.g., spreading seeds), its loss could cause a sudden regime shift ([Folke et al., 2004](#)).

Some commentators are sanguine about the possibility of disrupting ecological functions, because one likely result is a reduction of resources available to sustain wild animal populations.³ On their view, doing away with predation and high birth rates is a win-win relative to inaction: Either the ecosystem remains fertile but fundamental causes of suffering are eliminated, or it can support the birth of fewer unhappy wild animals ([Faria, 2023](#), p. 96; [Johanssen, 2020](#), p. 41). On the other hand, species with fast life history strategies may better tolerate habitat destruction ([Keesing & Ostfeld, 2021](#), p. 6), leading to an increased proportion of individuals with fast life history strategies, and perhaps even an increase in the overall number of wild animals.

Overall, disrupted ecosystems can make a habitat more hostile to life in an absolute sense—there are fewer resources to go around. However, they are less hostile to life in a relative sense—some abundant species may displace less abundant ones. How to weigh up the welfare gain from the reduction in resources from the welfare loss due to a change in the composition of animals is unclear.

³ To distinguish complex cluelessness due to ecosystem change from complex cluelessness due to measuring aggregate welfare, this subsection does not question the assumption that wild animal welfare is negative on net in the aggregate.



Abundance Versus Sentience

A concern with aggregate welfare would suggest prioritizing the most abundant groups of animals. On the other hand, only sentient animals have the capacity for welfare, so aggregate welfare would also prioritize the animals that have the highest likelihood of sentience. Among broad groups of wild animals, there is a negative association between abundance and probability of sentience ([Duffy, 2023b](#))⁴. The same criterion seems to provide strong reasons for prioritizing different groups of animals.

Compounding the complexity, estimating the probability of sentience is fraught. For many species, especially those with relatively low prior probabilities of sentience, scientists have never measured putative indicators of sentience ([Crump et al., 2022](#); [Gibbons et al., 2022](#)). In cases where they have, there is uncertainty about the internal validity and replicability of the results ([Diggles et al., 2023](#)). In cases where they have not, it is unclear what to assume in the meantime. Is absence of evidence some evidence of absence, because scientists are good at inferring when testing for an indicator of sentience is not worthwhile? Or should you expect that tests of sentience would come up positive because scientists underestimate the complexity of species who are distantly related to humans? More fundamentally, there is no consensus about the relevance of different indicators. Different approaches to sentience that are each plausible reach opposing verdicts for what types of evidence are central ([Schwitzgebel, 2020](#)).

Confusion about sentience introduces complex cluelessness to many wild animal welfare interventions. For example, installing bird-safe glass onto buildings could prevent the untimely death of over a billion birds each year ([Klem Jr., 2015](#)). Many birds consume invertebrates as part of their diet; should an evaluation of how bird-safe glass affects animal welfare account for indirect effects on prey invertebrate species? On the one hand, if invertebrates are not sentient, then billions of birds are helped and nobody is hurt.⁵ Moreover, it is not clear how to assess the indirect effects on invertebrates: How much longer would they live, and how well-off are they during that additional time? On the other

⁴ The negative correlation is far from perfect, with fruit flies and honeybees receiving similar probabilities of sentience as carp and salmon. This is likely because [Duffy's \(2023b\)](#) model treats absence of evidence as evidence of absence (p. 3), and there is a lot of research on fruit flies and honeybees.

⁵ This analysis ignores the fact that birds are interdependent with other vertebrate and non-prey species as well. It also assumes that the intervention is net-positive for the birds themselves, even though one could express doubt about how much better off they are living longer and enduring death by some other means. But restricting the discussion to just prey invertebrates is sufficient to make the case for complex cluelessness.

hand, the affected invertebrates likely far outnumber the affected birds, so if they are sentient then the effect of the intervention on invertebrates would dominate the overall effect on aggregate welfare. Furthermore, the likelihood that the invertebrates are exactly as well off regardless of whether their predators die from crashing into buildings in large numbers seems low, even though it is hard to justify any particular prediction. Overall, there are compelling reasons both to incorporate the effects of bird-safe glass on invertebrates and to ignore them.

Risk Aversion Helps Cope with Uncertainty

Complex cluelessness "feels particularly galling," and can result in inaction ([Greaves, 2016](#), p. 323). Inaction is itself a choice, and one that may be difficult to justify given the scale and urgency of wild animal welfare. With regards to scale, wild animals are the largest group of sentient beings, there may not be opportunities to do a comparable amount of good with lower risk.⁶

With regards to urgency, [Faria \(2023, pgs. 92-96\)](#) argues that inaction is only rational under two conditions. First, if the status quo is already the optimal scenario, then intervention can only reduce aggregate welfare. High juvenile mortality rates among wild animals alone renders the optimality of the status quo implausible. Second, even if the status quo is bad, inaction is still justified if all available actions have negative expected value. So long as there is at least one action with positive expected utility, action is required. Any benefits of further reducing risk beyond what is required to attain positive expected utility must be weighed against how long such efforts would take, since waiting itself locks in negative outcomes:

Though we expect doctors to be cautious when deciding upon and administering a treatment, we also think that the appropriate degree of caution is sensitive to how urgent the situation is. In a medical emergency where the patient will soon die without intervention, we would normally say that fast action on the doctor's part is morally required. But if that's the conclusion we draw in the medical case, we should draw a similar conclusion about intervention in the wild because waiting to intervene has significant negative consequences. Waiting means that more generations of r-Strategists will come into existence and live terrible lives, generations that we might have prevented from being born had we acted more quickly. Though we should wait until we have developed a degree of competence before

⁶ On the other hand, digital minds could eventually become the most abundant group of sentient individuals after the advent of digital sentience. If there will be less risky ways to improve the welfare of digital minds, then investing resources until digital minds arrive may have higher expected value than spending those resources on wild animals ([Šimčikas, 2023](#)).

intervening (much the way we expect those who respond to medical emergencies to undergo medical training), we need not wait until we have achieved perfect reliability. ([Johanssen, 2020](#), p. 42)

However, using standard expected value maximization to make decisions in the face of complex cluelessness may make stakeholders uncomfortable. To see why, imagine an intervention that will either increase or decrease aggregate welfare by 100 units. Due to complex cluelessness, there is uncertainty about the probability of increasing aggregate welfare. The best you can do is a uniform distribution, bounded at .35 and .67. Standard expected value maximization would simply take a weighted average of the probabilities ([Carr, 2020](#), p. 2744). Implementing the intervention would maximize expected value because $(.35 + .67/2) = .51$. Yet, it feels brazen to implement an intervention when the likelihood of doing good is only slightly higher than the likelihood of making the world worse. Standard expected value maximization is "risk-neutral," giving no special consideration to the distinct possibility of doing harm. It also feels unduly precise, in that it draws no distinction between whether the probability of success was .51 to begin with or only became .51 after averaging across a wide distribution. Standard expected value maximization is "ambiguity neutral," giving no special consideration to the degree of spread around probability estimates.

[Greaves \(2016, pgs. 327-329\)](#) considers the possibility that modeling complex cluelessness requires *imprecise credences*—that is, when there are multiple plausible probability assignments, they cannot simply be averaged together. When different probability assignments *taken individually* imply that different actions maximize expected value, it is indeterminate which action expected value maximization recommends. Although imprecise credences explicitly acknowledge ambiguity, it is unclear how to choose among potential actions when different probability assignments make opposing recommendations. In the present case, elements below .50 would recommend against the intervention; elements above .50 would recommend the intervention, and .50 would be indifferent between intervention and non-intervention. One candidate for how to make decisions in these situations, the *maximality* rule, claims that it is rationally permissible to choose among any of the indeterminate actions ([Mogensen, 2021](#)).

For many large-scale wild animal welfare interventions, it is difficult to rule out either probability assignments relative to which expected values are negative or probability assignments relative to which expected values are positive. The maximality rule therefore can make inaction in wild animal welfare defensible. That said, inaction is neither more nor less (or even equally!) preferable to action. As a result, the maximality rule cannot explain

or justify any systematic tendencies in how you choose among permissible actions. For instance, you might prefer to implement an intervention when the probability of a net-positive outcome is a uniform distribution between .49 and 1, even though you prefer inaction when the distribution is .35 to .67. You might also want an alternative to risk neutrality when evaluating an action according to any particular probability function.

Of course, stakeholders do not always know in advance which actions they prefer based on the available evidence. In the absence of guidance tailored to their risk attitudes, they may err towards inaction when they are unsure what to do. Explicitly incorporating aversion to risk and ambiguity into decision-making can help make preferences consistent across situations where it is clear what to do and situations where it is unclear what to do. As the following sections show, risk aversion often does recommend decisive action.

Three Types of Risk Aversion

"Risk aversion" is a broad term that encompasses distinct attitudes. This report overviews three types.⁷ This section begins by introducing an approach to wild animal welfare that is compatible with all three types of risk aversion but incompatible with risk neutrality.

Many wild animal advocates prefer interventions that do not change the number or composition of wild animals. [Šimčikas \(2022\)](#) argues that reducing anthropogenic ocean noise could be a cost-effective wild animal welfare intervention because noise may stress out a large number of marine animals. However, his enthusiasm is tempered by the possibility that indirect effects swamp the calculus:

Aquatic noise disturbs marine animals and affects ecosystems at many levels...Other than the first effect (direct stress), it's unclear whether these effects are overall positive or negative for animal welfare. I have encountered this same problem when analyzing most WAW [wild animal welfare] interventions. That is, I found that their most important effects are on the numbers of various wild animal species and it's very difficult to determine whether these effects are good or bad for WAW. Reducing aquatic noise seems a bit more robustly good for WAW than many other WAW interventions, partly because it seems at least plausible that the stress it causes directly could outweigh these indirect effects on animal numbers.

One way to avert complex cluelessness would be to leave current population dynamics intact. A paradigmatic example is "humane pesticide" for insect "pest" populations. Relative to the most popular pesticides, a humane pesticide would either kill insects faster, induce

⁷The exposition here is brief and informal, which is sufficient for the qualitative nature of the discussion. Formal overviews are available in [Clatterbuck, 2023](#) and [Duffy, 2023c](#).

less severe pain, or both. Intuitively, a concern that pesticide causes insects pain would motivate less pesticide use. However, less pesticide use might increase pest insect populations. Because it is unclear "whether insect lives are worth living and whether we should kill more or fewer of them," ([Tomasik, 2007](#)) killing the same number of them more humanely represents a large welfare gain with little downside risk. Similarly, [Howe \(2020\)](#) draws attention to the unclear effects pesticide reduction would have on interdependent species:

One of the biggest challenges to most wild animal welfare interventions is figuring out the net effect on all the species in the community. This is especially true when population sizes change. For example, falling wolf numbers might lead to rising coyote numbers, leading to falling fox numbers, leading to rising mouse numbers, leading to rising tick numbers—or they might not. If we can replace one insecticide with another that kills the same number of insects less painfully, we can reduce suffering without changing the total insect population. That significantly narrows the range of possible unintended consequences.

The same considerations apply to other areas of pest control, not just insects. For instance, most anti-rodenticide advocates push for alternatives that would reduce the birth rate of pest rodent populations, such as removing accessible sources of food and water ([Elmore et al., 2023](#)). But given uncertainty about how large a rodent population should be for optimal welfare, you could instead advocate for maintaining current population levels and killing them with traps instead of rodenticide. Although traps are painful if they do not cleanly maim rodents, in aggregate they probably cause far less harm than rodenticide.

Risk aversion is not a free lunch. Beyond pest management, it is unclear how to reduce suffering without affecting the abundance or composition of wild animals. Moreover, the upside of the remaining interventions is limited by leaving the basic situation in place. Pest populations do not get to be free from suffering during their daily life or live any longer, they only get to suffer less during death.

Outcome Risk Aversion

Outcome risk aversion recommends avoiding actions that risk the worst-case outcomes. This preference is represented by modifying the expected value of an action's possible outcomes such that the worse outcomes receive more decision weight. *Risk-weighted* expected utility orders outcomes (*not* discounted by probability) from worst to best. It then treats the worst outcome as a certainty, and discounts the improvement of each other outcome on the next-worst outcome by a function, r , of p , its probability ([Buchak, 2023](#)).



Consider how an outcome risk-averse actor who sets $r = p^2$ compares to a risk-neutral actor in deciding whether to swap out inhumane pesticide for humane pesticide or instead eliminate pesticide use altogether (see Table 1). There is certainty about the effects of each alternative on the badness of death and population size, but there is uncertainty about the welfare of pest insects before the experience of dying. There is a .2 probability of a "Hard Life"– life before death on the net has a utility of -50 for each insect– and a .8 probability of a "Neutral Life"– life before death on the net has a utility of 0.

Intervention	Neutral Life $p = .8$	Hard Life $p = .2$
Inhumane Pesticide (status quo)	$N = 100$ Utility Before Death = 0 Utility During Death = -20	$N = 100$ Utility Before Death = -50 Utility During Death = -20
Humane Pesticide	$N = 100$ Utility Before Death = 0 Utility During Death = -10	$N = 100$ Utility Before Death = -50 Utility During Death = -10
Eliminate Pesticide	$N = 200$ Utility Before Death = 0 Utility During Death = 0	$N = 200$ Utility Before Death = -50 Utility During Death = 0

Table 1. Factors affecting the expected value of pesticide interventions

In the status quo, there are 100 insects per generation. They all die by inhumane pesticide, an event which has a utility of -20. The risk-neutral expected utility is $.8*(100*-20) + .2*((100*-20) + (100*-50)) = -3000$. The risk-weighted expected utility is $-7000 + .64*(5000) = -3800$.

Humane pesticide halves the disutility of death and has no other welfare-relevant consequences. The risk-neutral expected utility is $.8*(100*-10) + .2*((100*-10) + (100*-50)) = -2000$, a 1000-unit improvement on inhumane pesticide. The risk-weighted expected utility is $-6000 + .64*(5000) = -2800$, an 1000-unit improvement on inhumane pesticide. The favorability of humane pesticide relative to inhumane pesticide does not change across risk-neutral expected utility and risk-weighted expected utility because humane pesticide does nothing about the possibility that life before death is net-negative.

Eliminating pesticide use reduces the disutility of death to 0 (e.g., all insects instantaneously die from predation instead), but it doubles the population of insects. The risk-neutral

expected utility is $8*(200*0) + .2*((200*0) + (200*-50)) = -2000$, a 1000-unit improvement on inhumane pesticide and 0-unit improvement on humane pesticide. Risk-neutrality is indifferent between humane pesticide and eliminating pesticide use. The risk-weighted expected utility is $-10000 + .64*(200*0) = -10000$, a 6200-unit reduction from inhumane pesticide and 7200-unit reduction from humane pesticide. Outcome risk aversion heavily punishes the elimination of pesticide for exacerbating the worst-case outcome. Humane pesticide cannot exacerbate the worst-case outcome, so it is the uniquely rational decision under outcome risk aversion.

Difference-Making Risk Aversion

Difference-making risk aversion gives special consideration to the likelihood that the action will be an improvement on not having acted at all. [Clatterbuck \(2023\)](#) and [Duffy \(2023c\)](#) modify the risk-weighted expected utility function to model *difference-making risk-averse expected utility*. The outcome that makes the least difference (or the largest negative difference) is treated as a certainty. The function m modifies the other outcomes by their probability of making a difference relative to the status quo.

Set m to $= p^2$. Because inhumane pesticide is the status quo, both its risk-neutral difference-making expected utility and difference-making risk averse expected utility are 0. For humane pesticide, there is a certainty of improving on inhumane pesticide by 1000 units. For eliminating pesticide, the difference-making risk averse expected utility is $-3000 + (.64*2000) = -1720$, 2720 units worse than humane pesticide. Eliminating pesticide gets some credit here for the high likelihood of making death completely painless. The penalty it gets for increasing the insect population is more modest than it was for outcome risk aversion because the status quo was already so bad on the assumption life before death is net-negative. Humane pesticide again prevails, because there is a guarantee of improving on the status quo.

Ambiguity Aversion

Ambiguity aversion penalizes ignorance rather than risk per se. Whereas risk denotes the probability of undesirable outcomes, ignorance refers to a lack of clarity about what probabilities to assign to different outcomes ([Peterson, 2017](#), ch. 2-3). If there is exactly a 75% chance of a negative outcome, then there is high risk but no ignorance. If the chance of a negative outcome is somewhere between .0001% and 1%, then there is low risk but substantial ignorance. There are many ways of modeling ambiguity aversion ([Buchak, 2023](#), pgs. 181-183). For simplicity, consider α -maximin, which assigns weight α to the



upper-bound of the distribution and $1-\alpha$ to the lower-bound. An actor is ambiguity-averse when the value of α is lower than .5.

This example will use $\alpha = .25$. Instead of a .8 probability that life before death has a net utility of 0 rather than 5, there is now a uniform probability distribution ranging from .7 to .9. For inhumane pesticide, the upper-bound expected utility is $.9*(100*-20) + .1*((100*-20) + (100*-50)) = -2500$ and the lower-bound expected utility is $.7*(100*-20) + .3*((100*-20) + (100*-50)) = -3500$. The α -maximin expected utility is $(.25*-2500) + (.75*-3500) = -3250$.

For humane pesticide, the upper-bound expected utility is $.9*(100*-10) + .1*((100*-10) + (100*-50)) = -1500$ and the lower-bound expected utility is $.7*(100*-10) + .3*((100*-10) + (100*-50)) = -2500$. The α -maximin expected utility is $.25*(-1500) + .75*(-2500) = -2250$. For eliminating pesticide, the upper-bound expected utility is $.9*(200*0) + .1*((200*0) + (200*-50)) = -1000$ and the lower-bound expected utility is $.7*(200*0) + .3*((200*0) + (200*-50)) = -3000$. The α -maximin expected utility is $.25*(-1000) + .75*(-3000) = -2500$.

Humane pesticide is the uniquely rational choice again, because its efficacy is not tied to whether life before death is bad or not. In contrast, eliminating pesticide gets penalized because of the likelihood of aggravating the aggregate disutility of life before death. However, the penalty is much softer than it was for outcome risk aversion, so it is still more choiceworthy than inhumane pesticide.

The Value– and Rarity– of Robustness

At least in the toy example explored here, humane insecticide is favored by all three types of risk aversion. Call this property *risk-aversion robustness*. This section argues that, all else equal, robustness is desirable. However, part of the reason that robustness is indicative of choiceworthiness is because, in general, different forms of risk aversion disagree in their recommendations. Accordingly, this section also demonstrates a lack of robustness for many areas of wild animal welfare.

Convergent Evidence

Catering to the risk attitudes of stakeholders provides a pragmatic reason to incorporate risk aversion into expected value calculations. However, it still matters whether risk aversion is in fact rational. Even if some version of risk aversion is valid, it does not follow that all types are ([Greaves et al., 2022](#)). Until there is greater clarity on how to adjudicate



between different decision-making criteria, you can at least assess the sensitivity of results to different types of risk aversion. If the probability that different types of risk aversion are valid are at least partly independent of each other, robustness across multiple types of uncertainty provides convergent evidence that a cause or intervention is choiceworthy.

Even more robust would be interventions that are favored both on risk-neutral grounds and on risk-averse grounds. When the degree of risk aversion is weak, robustness is not as diagnostic of choiceworthiness because risk-aversion and risk-neutrality are highly positively correlated. So, it is more helpful to look for interventions that gain the approval of both risk-neutral and strongly risk-averse actors, but by definition these cases should be rare. It is easier to identify considerations *against* an intervention that resonate on both risk-neutral and risk-averse grounds. In wild animal welfare, a fear of alienating others can be justified by many risk attitudes. From a risk-neutral perspective, an intervention may appear to have high expected value because it is technically feasible and the foreseeable side effects have been planned for. Nevertheless,

explicitly stewarding nature for specific purposes might be perceived negatively by the public...In other words, certain interventions might superficially seem very cost-effective, but only in the absence of political and coordination costs. If public opposition to these interventions are large enough, the cost-effectiveness of working towards implementing them will be low. ([Liedholm, 2019, p. 10](#)).

The largest reputational risk comes not from jeopardizing the actual project under consideration, but rather from tainting the wild animal welfare movement for years to come. [Tomasik \(2015\)](#) warns, "before we become too gung ho about eliminating natural ecosystems, we should also remember that many other humans value wilderness, and it's good to avoid making enemies or tarnishing the suffering-reduction cause by pitting it in direct opposition to other things people care about" (p. 146). That is, the expected value of the cause area as a whole would shrink if the tractability of future interventions decreased as a result of pushing for radical changes to ecosystems too quickly. Outcome and difference-making risk-averse actors will agree that public relations concerns mitigate against controversial interventions, because the wild animal welfare movement may be the only way to stop widespread suffering in natural environments. Ambiguity aversion also tends to oppose controversial interventions, as uncontroversial interventions need not account for the possibility of a backlash of unclear magnitude.



Divergent Recommendations

If different forms of risk aversion always agreed in their recommendations, there would be little value in distinguishing between them. This section uses examples from wild animal welfare to lay out two situations where different types of risk aversion reliably give different advice.

Probability Versus Value

When the alternatives under consideration vary in whether they (a) have a lower probability of success but higher value conditional upon success or (b) have a higher probability of success but lower value conditional upon success, outcome and difference-making risk aversion tend to reach opposite verdicts. A clear example is how to make trade-offs between the abundance of wild animals and their probability of sentience (see [Clatterbuck, 2023](#) for a toy example). Outcome risk aversion would *magnify* the priority of protecting the most abundant group of animals, because ignoring their potential needs poses the largest risk of unnecessary harm. Difference-making risk aversion instead stresses the probability of helping some animals, and therefore prioritizes groups of more certain sentience.⁸ For ambiguity aversion, species that have vague probabilities of sentience are penalized relative to those who have well-established probabilities. In theory, a species with a precise but lower probability of sentience could be favored over a species with a higher but fuzzier probability of sentience.

Benefits of Research

Ambiguity aversion also introduces a new alternative—conducting research to reduce uncertainty about the expected value of each intervention option. For example, you could conduct sentience research on abundant species for which there is little or no sentience

⁸ As [Clatterbuck \(2023\)](#) points out, however, you can increase the overall probability of doing some good by helping multiple groups of animals who have relatively independent probabilities of sentience. How independent these probabilities are is difficult to judge—sentience may be a highly conserved trait, or its adaptive value might depend on occupying specific evolutionary niches. One plausible case is that the probability of sentience among vertebrates and invertebrates is fairly independent. Vertebral columns probably evolved around 560-520 million years ago, a period when it is unclear whether sentience existed yet. [Feinberg and Mallatt \(2013\)](#) argue that some of the earliest vertebrates possessed the capacity for subjective experience, while invertebrates did not yet. If so, then whether abundant groups of invertebrates evolved sentience later is at least partially independent of the fact that vertebrates apparently already had it. Difference-making risk aversion, then, is more tolerant of resource allocation towards, say, insects and fishes rather than fishes alone, even though fishes generally have higher probabilities of sentience than insects ([Duffy, 2023b](#), Table 1).

research. For outcome risk aversion, it is already clear that worst-case outcomes involve failing to help highly abundant groups of animals. Because sentience research is unlikely to deliver knock-down evidence against the sentience of these groups, more accurate probabilities of sentience is unlikely to affect decision-making.

Over the long run, sentience research would increase the share of projects that make at least some positive difference. Difference-making risk aversion could regard research as an investment in the probability of success of future interventions. Yet, the probability that any one project on sentience will make a material difference to wild animals is low. This is especially so if the audience for the research is also difference-making risk averse— their behavior would not change if the research decreased the probability of sentience among abundant populations, only if it increased it.⁹

There are two reasons why difference-making risk aversion could favor funding sentience research after all. First, even though any one research project is unlikely to make a difference, there may be occasional breakthroughs that substantially change the probability that a group of animals is sentient. Conducting a large number of projects would increase the probability that at least one of them results in a breakthrough. Similarly, it could be that no one project can make a difference on its own, but funding an entire body of evidence can.

Second, to improve on doing nothing, you have to avoid crowding out other research funders. [Trammell \(2021, pgs. 30-31\)](#) claims that philanthropists disfavor basic research because they neglect benefits that disproportionately accrue in the future. Funding sentience research now should speed up the pace at which it is conducted rather than merely displace other funding. However, crowding out is a general argument for funding research that has no obvious immediate application. The short-term utility of sentience research may actually be more apparent to funders than other areas of wild animal welfare science, such as how various ecosystem dynamics affect welfare.

Risk Aversion Affects Animal Welfare Cause Prioritization

So far, the analysis considers how risk aversion affects the choice among wild animal welfare interventions. In reality, risk aversion might favor allocating resources to an entirely different cause area. For those working in animal welfare, farmed animal welfare would be the natural alternative, due to the large number of farmed animals relative to

⁹ Thanks to Hayley Clatterbuck for this point.



other types of captive animals ([Animal Charity Evaluators, 2016](#)). Possibly, the reason that farmed animal welfare receives more funding than wild animal welfare ([Animal Charity Evaluators, 2023](#)) is that stakeholders believe that farmed animal interventions are less risky. As it turns out, risk aversion does not robustly favor either cause area.

The most relevant variables for outcome risk aversion are the number of animals at stake and their welfare. Wild animal welfare clearly surpasses farmed animal welfare in the number of animals involved. In terms of how poorly animals are doing, there is a high probability that farmed animals suffer a lot in general. However, many believe that wild animals also generally suffer a lot. For them, wild animal welfare is at least comparably urgent to farmed animal welfare.

There is greater uncertainty about the overall situation of wild animals, though. Ambiguity aversion therefore prefers helping farmed over wild animals because it is clearer how much help farmed animals need. Of course, ambiguity aversion could also recommend research that would help both fields, such as investigating the sentience or welfare needs of insects. As discussed in the previous section, difference-making risk aversion could also favor a portfolio of research. Because it is easier to translate research on farmed animals into welfare reforms, difference-making risk aversion would probably favor research on intensively farmed species over especially abundant wild species. On the other hand, welfare research on abundant wild species may be particularly neglected, and thus more likely to make a counterfactual difference.

Conditional upon implementing large-scale interventions in the near-term, ambiguity aversion and difference-making risk aversion both favor farmed animal welfare.¹⁰ Research has identified interventions that have a non-trivial probability of helping a large number of farmed animals (e.g., [Duffy, 2023a](#); [Jalil et al., 2023](#); [Šimčikas, 2019](#)). There is no comparable evidence base for wild animal welfare interventions, so the probability of success is lower and more uncertain. That said, difference-making risk aversion would also favor diversifying investments across farmed and wild animal interventions insofar as their

¹⁰ For those planning over a longer time horizon, difference-making risk aversion and ambiguity aversion could both recommend investing resources for now, and waiting for particularly tractable, high-impact opportunities in either farmed or wild animal welfare. However, difference-making risk aversion does need to account for the fact that waiting to allocate assets increases the probability of losing control over them ([Trammell, 2021](#)).



probabilities of success are independent.¹¹ To maximize the likelihood of doing some good in wild animal welfare, it is advisable to focus on less ambitious interventions. To illustrate, [Elmore et al. \(2023\)](#) found that state-level rodenticide bans are unpopular with the U.S. public. However, several reforms that did not ban rodenticides but would still modestly reduce their use were popular. If the reforms go well, then bans could become more popular over time as people come to believe that rodenticides are not as necessary as they had thought.

Conclusion

This report introduced various types of risk aversion on the premise that complex cluelessness is responsible for the relatively modest allocation of resources to wild animal welfare in spite of its unparalleled scale. By providing principled tools to incorporate uncertainty into expected value calculations, stakeholders may be more comfortable taking action when they know the status quo is suboptimal. Rethink Priorities recently published a cross-cause effectiveness model that enables users to model different types of risk aversion when comparing cause areas ([Shiller et al., 2023](#)). The model has not yet been applied to wild animal welfare, in large part because estimates do not yet exist for how much wild animals can be helped per some dollar amount. As the field begins developing direct interventions, the cross-cause effectiveness model can help determine whether wild animal welfare is cost-competitive with fields like farmed animal welfare.

In the meantime, stakeholders can at least reflect on the direction that different types of risk aversion push expected value calculations. Outcome risk aversion favors wild over farmed animal welfare, and particularly favors abundant species, even if their probability of sentience is low. In contrast, difference-making risk aversion mainly prioritizes farmed animal welfare and species with higher probabilities of sentience. It also encourages diversification across uncorrelated interventions and neglected interventions. Ambiguity aversion prioritizes evidence-based farmed animal interventions and research on areas of high uncertainty, like the sentience of understudied groups of animals. All three types favor interventions that reduce suffering without changing the population dynamics of wild animals. And they all agree with a risk-neutral perspective on avoiding controversial interventions.

¹¹ One reason for assuming partial independence is that the fields differ in ways that most people regard as morally relevant. In particular, how people treat farmed animals is about harms of commission, whereas indifference to non-anthropogenic sources of suffering is a sin of omission. That said, it is beyond the present scope to demonstrate that probabilities of success vary more between the two cause areas than they do within them.

Admittedly, the hypothesis that unmodeled risk aversion is responsible for the neglectedness of wild animal welfare is highly speculative. Appealing to complex cluelessness may prove too much, as it occurs in many other contexts ([Greaves, 2016](#), p. 334). Complex cluelessness alone is not sufficient to explain decision paralysis; a more unique aspect of wild animal welfare must also play some role. Possibly, wild animal welfare stands out for its unusually high stakes¹², which makes the possibility of inadvertently doing harm particularly salient.

Alternatively, wariness about intervention into natural habitats could be due to uncertainty about its *moral* permissibility. Some non-consequentialist philosophers argue that humans should not radically change natural habitats or their inhabitants to increase aggregate welfare ([Donaldson & Klymycka, 2013](#), p. 155; [Korsgaard, 2018](#), pgs. 186-187; [Nussbaum, 2023](#), p. 119). If their intuitions are widespread, then the average animal advocate may have reservations about large-scale interventions. However, these philosophers do not counsel inaction, but instead encourage a proactive reduction of human interference into natural habitats, except in certain limited circumstances. Inaction due to moral uncertainty would again suggest a misalignment between the calculus decision-makers profess using—risk-neutral expected value maximization— and the range of considerations that they believe are relevant— moral views that reject aggregate welfare as a meaningful concept.

Moral uncertainty analysis can recommend compromises between moral theories with opposing prescriptions. However, inaction is an unlikely compromise between theories that (might) demand habitat destruction and theories that demand habitat preservation, as they all regard the status quo as undesirable. Although a complete analysis is beyond scope, certain types of meat reduction interventions are plausible candidates. Some types of farming, especially raising cattle, are responsible for large amounts of deforestation ([Ritchie, 2021](#)), which can affect the number and composition of wild animals via regime shifts and changes in net primary productivity ([Matheny & Chan, 2005](#); [Smil, 2013](#)). A utilitarian perspective would not prioritize reducing beef consumption because there are relatively few livestock and they generally have higher welfare than other farmed animals ([Mathur, 2022](#)). Moreover, if wild animals have bad lives, then the conversion of forests into less productive grasslands could increase aggregate welfare. A compromise between the two worldviews might be focusing on reducing consumption of numerous, small-bodied animals, who generally fare worse than livestock and occupy less land. This approach

¹² Stakes-sensitive actors may consider weighted-linear utility theory as superior to risk-expected utility theory for modeling outcome risk aversion ([Bottomley & Williamson, 2023](#)).

would strongly reduce farmed animal suffering (very good on both worldviews) while also modestly reducing land use (somewhat good on the non-consequentialist view, and potentially somewhat bad on the utilitarian view).

Notably, risk aversion may also be appropriate to factor into moral uncertainty analysis. [MacAskill et al. \(2020\)](#), p. 48) argue that empirical and moral uncertainty should be handled the same way with respect to risk attitudes– if risk-neutrality is appropriate for one, it is appropriate for the other, but if risk aversion is appropriate for one then it is for the other too. Identifying which theories of animal ethics receive the most decision weight under each type of risk aversion would be a valuable future direction.

Acknowledgments



The post was written by William McAuliffe. Thanks to Bob Fischer, Hayley Clatterbuck, Neil Dullaghan, and Daniela Waldhorn for helpful feedback. The post is a project of [Rethink Priorities](#), a global priority think-and-do tank, aiming to do good at scale. We research and implement pressing opportunities to make the world better. We act upon these opportunities by developing and implementing strategies, projects, and solutions to key issues. We do this work in close partnership with foundations and impact-focused non-profits or other entities. If you're interested in Rethink Priorities' work, please consider subscribing to [our newsletter](#). You can explore our completed public work [here](#).

References

- Animal Charity Evaluators. (2016). Why farmed animals? Retrieved November 10, 2023, from <https://animalcharityevaluators.org/donation-advice/why-farmed-animals/>
- Animal Charity Evaluators. (2023). Why wild animals? *Animal Charity Evaluators*. Retrieved November 10, 2023, from <https://animalcharityevaluators.org/donation-advice/why-wild-animals/>
- Animal Ethics. (n.d.). Helping animals in the wild. Retrieved November 10, 2023, from <https://www.animal-ethics.org/helping-animals-in-the-wild/>



- Animal Ethics. (n.d.). Strategic considerations for effective wild animal suffering work. Retrieved November 10, 2023 from <https://www.animal-ethics.org/strategic-considerations-for-effective-wild-animal-suffering-work/>
- Bar-On, Y. M., Phillips, R., & Milo, R. (2018). The biomass distribution on Earth. *Proceedings of the National Academy of Sciences*, 115(25), 6506–6511. <https://doi.org/10.1073/pnas.1711842115>
- Bottomley, C., & Williamson, T. L. (2023). Rational risk-aversion: Good things come to those who weight. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.13006>
- Browning, H., & Veit, W. (2023). Positive wild animal welfare. *Biology & Philosophy*, 38(2), 14. <https://doi.org/10.1007/s10539-023-09901-5>
- Buchak, L. (2023). How should risk and ambiguity affect our charitable giving? *Utilitas*, 35(3), 175–197. <https://doi.org/10.1017/S0953820823000055>
- Carr, J. R. (2020). Imprecise evidence without imprecise credences. *Philosophical Studies*, 177(9), 2735–2758. <https://doi.org/10.1007/s11098-019-01336-7>
- Clatterbuck, H. (2023). The Risks and Rewards of Prioritizing Animals of Uncertain Sentience. Rethink Priorities. Retrieved November 10, 2023, from <https://rethinkpriorities.org/publications/the-risks-and-rewards-of-prioritizing-animals-of-uncertain-sentience>
- Crump, A., Browning, H., Schnell, A., Burn, C., & Birch, J. (2022). Sentience in decapod crustaceans: A general framework and review of the evidence. *Animal Sentience*, 7(32). <https://doi.org/10.51291/2377-7478.1691>
- Cuddington, K. (2019). Insect herbivores, life history and wild animal welfare. Rethink Priorities. Retrieved November 10, 2023, from <https://rethinkpriorities.org/publications/insect-herbivores-life-history-and-wild-animal-welfare>
- Delon, N., & Purves, D. (2018). Wild animal suffering is intractable. *Journal of Agricultural and Environmental Ethics*, 31(2), 239–260. <https://doi.org/10.1007/s10806-018-9722-y>
- Diggles, B. K., Arlinghaus, R., Browman, H. I., Cooke, S. J., Cooper, R. L., Cowx, I. G., Derby, C. D., Derbyshire, S. W., Hart, P. J., Jones, B., Kasumyan, A. O., Key, B., Pepperell, J. G., Rogers, D. C., Rose, J. D., Schwab, A., Skiftesvik, A. B., Stevens, D., Shields, J. D., & Watson, C. (2023). Reasons to be skeptical about sentience and pain in fishes and aquatic invertebrates. *Reviews in Fisheries Science & Aquaculture*, 0(0), 1–24. <https://doi.org/10.1080/23308249.2023.2257802>
- Donaldson, S., & Kymlicka, W. (2013). A defense of animal citizens and sovereigns. *Law, Ethics and Philosophy*, 143–160.

- Duffy, L. (2023a). A cost-effectiveness analysis of historical farmed animal welfare ballot initiatives. Rethink Priorities. Retrieved November 10, 2023, from <https://rethinkpriorities.org/publications/a-cost-effectiveness-analysis-of-historical-farmed-animal-welfare-ballot-initiatives>
- Duffy, L. (2023b). Welfare Range and P(Sentience) Distributions. Retrieved November 10, 2023, from https://docs.google.com/document/d/1xUvMKRkEOJQcc6V7VJqcLLGAJ2SsdZno0jTIUb61D8k/edit?usp=sharing&usp=embed_facebook
- Duffy, L. (2023c). How Can Risk Aversion Affect Your Cause Prioritization? Google Docs. Retrieved November 10, 2023, from https://docs.google.com/document/d/1CZ5S-Eayxr64z5YADYR9M3P2WTP4u2Pgb4N-ynYbbMU/edit?usp=sharing&usp=embed_facebook
- Elmore, H., McAuliffe, W., & McKay, H. D. (2023). Paths to reducing rodenticide use in the U.S. OSF Preprints. <https://doi.org/10.31219/osf.io/4ha57>
- Faria, C. (2022). *Animal Ethics in the Wild: Wild Animal Suffering and Intervention in Nature*. Cambridge University Press. <https://doi.org/10.1017/9781009119948>
- Feinberg, T. E., & Mallatt, J. (2013). The evolutionary and genetic origins of consciousness in the Cambrian Period over 500 million years ago. *Frontiers in Psychology*, 4, 51583. <https://doi.org/10.3389/fpsyg.2013.00667>
- Folke, C., Carpenter, S., Walker, B., Scheffer, M., Elmqvist, T., Gunderson, L., & Holling, C. S. (2004). Regime shifts, resilience, and biodiversity in ecosystem management. *Annual Review of Ecology, Evolution, and Systematics*, 35(1), 557–581. <https://doi.org/10.1146/annurev.ecolsys.35.021103.105711>
- Gibbons, M., Crump, A., Barrett, M., Sarlak, S., Birch, J., & Chittka, L. (2022). Can insects feel pain? A review of the neural and behavioural evidence. *Advances in Insect Physiology* 63, 155–229. <https://doi.org/10.1016/bs.aip.2022.10.001>
- Greaves, H. (2016). XIV—Cluelessness. *Proceedings of the Aristotelian Society*, 116(3), 311–339. <https://doi.org/10.1093/arisoc/aow018>
- Greaves, H., Mogensen, A., MacAskill, W., & Thomas, T. (2022). On the desire to make a difference. *Global Priorities Institute Working Paper Series*. Retrieved November 10, 2023, from <https://globalprioritiesinstitute.org/on-the-desire-to-make-a-difference-hilary-greaves-william-macaskill-andreas-mogensen-and-teruji-thomas-global-priorities-institute-university-of-oxford/>
- Howe, H. (2020). Improving pest management for wild insect welfare. Wild Animal Initiative. Retrieved November 10, 2023, from <https://www.wildanimalinitiative.org/library/humane-insecticides>

- Jalil, A. J., Tasoff, J., & Bustamante, A. V. (2023). Low-cost climate-change informational intervention reduces meat consumption among students for 3 years. *Nature Food*, 4(3), Article 3. <https://doi.org/10.1038/s43016-023-00712-1>
- Johannsen, K. (2020). To assist or not to assist? Assessing the potential moral costs of humanitarian intervention in nature. *Environmental Values*, 29(1), 29–45. <https://doi.org/10.3197/096327119X15579936382644>
- Keesing, F., & Ostfeld, R. S. (2021). Impacts of biodiversity and biodiversity loss on zoonotic diseases. *Proceedings of the National Academy of Sciences*, 118(17), e2023540118.
- Klem, D. (2015). Bird–Window collisions: A critical animal welfare and conservation issue. *Journal of Applied Animal Welfare Science*, 18(sup1), S11–S17. <https://doi.org/10.1080/10888705.2015.1075832>
- Korsgaard, C. M. (2018). *Fellow creatures: Our obligations to the other animals*. Oxford University Press. <https://doi.org/10.1093/oso/9780198753858.003.0010>
- Liedholm, S. E. (2019). Long-term design considerations for wild animal welfare interventions. <https://www.wildanimalinitiative.org/blog/persistenceandreversibility>
- MacAskill, M., Bykvist, K., & Ord, T. (2020). *Moral Uncertainty*. Oxford University Press. <https://doi.org/10.1093/oso/9780198722274.001.0001>
- Matheny, G., & Chan, K. M. A. (2005). Human diets and animal welfare: The illogic of the ladder. *Journal of Agricultural and Environmental Ethics*, 18(6), 579–594. <https://doi.org/10.1007/s10806-005-1805-x>
- Mathur, M. B. (2022). Ethical drawbacks of sustainable meat choices. *Science*, 375(6587), 1362–1362. <https://doi.org/10.1126/science.abo2535>
- McMahan, J. (2010). The meat eaters. Opinionator. <https://archive.nytimes.com/opinionator.blogs.nytimes.com/2010/09/19/the-meat-eaters/>
- Mogensen, A. L. (2021). Maximal Cluelessness. *The Philosophical Quarterly*, 71(1), 141–162. <https://doi.org/10.1093/pq/pqaa021>
- Ng, Y.-K. (1995). Towards welfare biology: Evolutionary economics of animal consciousness and suffering. *Biology and Philosophy*, 10(3), 255–285. <https://doi.org/10.1007/BF00852469>
- Nussbaum, M.C. (2023). *Justice for Animals*. Simon and Schuster.
- Peterson, M. (2017, March 29). *An Introduction to Decision Theory*. Cambridge University Press. <https://doi.org/10.1017/9781316585061>
- Ritchie, H. (2021). Drivers of Deforestation. *Our World in Data*. <https://ourworldindata.org/drivers-of-deforestation>
- Schwitzgebel, E. (2020). Is there something it's like to be a garden snail? *Philosophical Topics*, 48(1), 39–64.
- Shiller, D., Bernard, D.S., Carter, C., Covarrubias, A., Davis, M.A., Dickens, M.,



- Duffy, L., & Wildeford, P. (2023). Rethink Priorities' cross-cause cost-effectiveness model: Introduction and overview. Retrieved November 10, 2023, from <https://forum.effectivealtruism.org/posts/pniDWyjc9vY5sjGre/rethink-priorities-cross-cause-cost-effectiveness-model>
- Šimčikas, S. (2019). Corporate campaigns affect 9 to 120 years of chicken life per dollar spent. Rethink Priorities. Retrieved November 10, 2023, from <https://rethinkpriorities.org/publications/corporate-campaigns-affect-9-to-120-years-of-chicken-life-per-dollar-spent>
- Šimčikas, S. (2022). Reducing aquatic noise as a wild animal welfare intervention. Rethink Priorities. Retrieved November 10, 2023, from <https://rethinkpriorities.org/publications/reducing-aquatic-noise>
- Šimčikas, S. (2023). Why I No Longer Prioritize Wild Animal Welfare. Retrieved November 10, 2023, from <https://forum.effectivealtruism.org/posts/saEQXBgzmdBob9GdH/why-i-no-longer-prioritize-wild-animal-welfare>
- Smil, V. (2013). *Harvesting the Biosphere*. MIT Press. <https://mitpress.mit.edu/9780262528276/harvesting-the-biosphere/>
- Tomasik, B. (2015). The importance of wild-animal suffering. *Relations. Beyond Anthropocentrism*, 3(2). <https://doi.org/10.7358/rela-2015-002-toma>
- Tomasik, B. (2007). Humane insecticides. Retrieved November 10, 2023, from https://reducing-suffering.org/humane-insecticides/#These_techniques_seem_net_good_if_not_replacing_insecticides
- Trammel, P. (2021). Patient Philanthropy in an Impatient World. Google Docs. Retrieved November 10, 2023, from https://docs.google.com/document/u/1/d/1NcfTgZsqT9k30ngeQbappYyn-UO4vltjkm64n4or5r4/edit?usp=embed_facebook
- Waldhorn, D. R., & Autric, E. (2022). Shrimp: The animals most commonly used and killed for food production. OSF Preprints. <https://doi.org/10.31219/osf.io/b8n3t>
- Wild Animal Initiative (2022). Annual report: Wild animal welfare science takes flight. Retrieved November 10, 2023, from <https://www.wildanimalinitiative.org/transparency>
- World Population Clock: 8.1 Billion People (LIVE, 2023) - Worldometer. (n.d.). Retrieved November 10, 2023, from <https://www.worldometers.info/world-population/>