The Intrinsic Valuation of Biodiversity Loss*

Adrien Bilal
Stanford University

Noémie Pinardon-Touati
Columbia University

Conor Walsh
Columbia University

July 2026

Abstract

We explore the welfare costs of the loss of animal life in a consequentialist total utilitarian framework that incorporates the intrinsic value of living beings through their sentience. Moral philosophy and neuroscience define sentience as the capacity for valenced experience. Observable performance on cognitive tests represents a conservative lower bound on sentience under plausible neuroscientific assumptions. We estimate the relationship between cognition and neuron counts to cardinalize sentience. We collect population and neuron data for a wide range of animal species, and aggregate them into a single welfare measure. Across a wide range of parameters, the aggregate sentience losses of animals in the past half century outstrip the sentience gains to humans from population growth by at least 10,000-fold. Breaking even in welfare terms requires assigning a correspondingly small welfare weight to animals. Including the benefits of economic growth to humans does little to change these overall welfare losses.

Keywords: Growth, Animal Welfare, Welfare Economics, Biodiversity Loss

JEL Codes: D61, D63, O44, Q57, Q56

*Acknowledgments: We are grateful to Sisi Qiu and Andres Garcia for excellent research assistance.

1. INTRODUCTION

Economic growth has transformed the lives of billions of humans. Steady improvement since the Industrial Revolution in frontier economies, and more recently in the developing world, has brought a total transformation of what it means to live and thrive on this planet. By any reasonable metric, a given human almost anywhere in the world lives a life better than their grandparents, and one unimaginable to their distant ancestors. As a result, economic growth has sustained continuous increases in the human population.

At the same time, the spread of humanity has been accompanied by dramatic declines in wild population counts in the natural world. Many species have been lost entirely through extinction, and many more of those that remain have suffered large population declines. These trends have occurred through a combination of habitat loss, competition for resources, and pollution, all direct consequences of economic growth.

This paper attempts to account for what has been lost during humanity's rise. To do so, we consider a consequentialist total utilitarian view, according to which there exists a welfare function aggregated across all relevant living beings. Moral philosophy argues that living beings that exhibit sentience may reasonably be included in the welfare function, where sentience is defined as the capacity for valenced experience. While sentience is difficult to measure, observable performance on cognitive tests represents a conservative lower bound on sentience under plausible neuroscientific assumptions. We estimate the relationship between cognition and neuron counts to cardinalize this lower bound on sentience. We collect population and neuron data for a wide range of animal species, and aggregate them into a single welfare measure. Across a wide range of parameters, the aggregate sentience losses of animals in the past half century outstrip the sentience gains to humans from population growth by more than 10,000-fold. Converting sentience to welfare and accounting for economic growth makes little difference to this comparison. Less conservative choices can imply relative losses in excess of 10^7. Equivalently, breaking even in welfare requires placing a relative welfare weight on animal sentience between 10^{-7} and 10^{-4}.

We clarify two aspects of our approach before describing it in more detail. First, we employ a capacious definition of biodiversity loss, which takes into account changes in overall wildlife populations. We focus primarily on changes in quantities and do not attempt to value species variety per se. Second, our approach is distinct from the instrumental valuation of wildlife such as ecosystem services. Our goal is to quantify

the intrinsic value of lost wildlife. We build on recent progress in moral philosophy and neuroscience to define which living beings to include in the welfare function, and what weight they might receive.

We start our analysis by defining the scope of the consequentialist total utilitarian approach. Only individuals with moral status enter the social welfare function. A broad consensus in moral philosophy, including utilitarianism, is that a living being has moral status if and only if it exhibits some degree of sentience. Sentience is defined as the capacity for valenced experience. Experience in this context means that there is "something it is like" for the living being to be in a given state, and can equivalently be defined as phenomenal consciousness. Valence is the property of experience being positive or negative for the living being. These properties encompass pain and pleasure, but also fear, anger, and their analogs in animals.

Sentience is a product of the brain: the vast majority of philosophers and neuroscientists hold that valenced experience is ultimately a physical phenomenon realized in the brain. In particular, neuroscientific evidence points to two brain regions that determine sentience: the midbrain (an evolutionarily more ancient brain region) and the pallium (an evolutionarily more recent brain region that also determines cognition). These brain regions are differentially developed across species. As a result, different species are likely to experience different degrees of sentience. For instance, vertebrates are generally accepted as exhibiting an ability for valenced experience. Given their distinct neurological architecture, insects and some families of molluscs are plausible candidates, but reasonable people may disagree about their moral status. The definition of sentience excludes plants, bacteria, and animals such as jellyfish, as they have either no neurological systems or extremely rudimentary ones, preventing them from being conscious.

Equipped with the notion of sentience that delineates which species might reasonably be included in the welfare function, we must then assign a cardinal utility to individuals of each species, sometimes called a welfare range in the context of animal welfare. Thus, we need a cardinal measure of sentience. Such a measure of sentience is still beyond existing neuroscientific measurement. However, neuroscientific evidence points to a plausible lower bound on the sentience of other animals relative to humans using measures of cognition.

Valence relies on two brain regions, the evolutionarily older midbrain and the evolutionarily more recent pallium, which controls cognition. Since the midbrain is

more similar across species than the pallium, the ability for valenced experience is likely more evenly spread than cognition across animals. In addition, neuroscientific functional arguments suggest that cognition scales more steeply in overall neuronal capacity than sentience. Therefore, variation in cognition across species is likely greater than variation in sentience, allowing us to construct a lower bound on the sentience of other animals relative to humans.

Direct measures of cognition are available for many species, by observing their ability to enumerate or solve puzzles to retrieve food. Leveraging again that cognition is tightly linked to neuronal capacity, we estimate the relationship between cognitive performance and total and pallial neuron counts for species for which both measures are available. We then predict cognition for all other species using their neuron counts. We obtain a cardinal measure of variation in cognition across all species with available neuron data, and thus a lower bound on the sentience of other animals relative to humans. This lower bound is cardinal under the important assumption that the cognition scale that we choose (test success rate relative to humans) represents cognition without the need for an additional transformation.

Of course, one can reasonably disagree with many of the steps involved in this calculation. While not devoid of challenges, this exercise has the advantage of being amenable to aggregation and quantitative analysis, and making underlying assumptions transparent. We thus view it as a useful first step in the quantification of the intrinsic value of biodiversity loss, but certainly not the last. To assess the role of each step, we report a wide range of sensitivity analyses.

Which species are the most sentient? Unsurprisingly, great apes are among the most sentient animals. Their sentience scores are often above 0.6 times that of humans. Sentience is correlated with body mass, but only imperfectly. In particular, many small animals are highly sentient. Mice have a sentience score of 0.06 relative to humans, while cows have a sentience score of 0.32, despite being 20,000 times heavier. Similarly, crows have a sentience score of 0.12, salmon of 0.003, and crickets of 0.002.

We then combine our measures of sentience with changes in wildlife population counts by species. We collect data on population trends at the species, order, and group levels for a wide range of animals. Population loss has been significant across all species. It ranges from around 45% for birds and mammals to 90% for herptiles (reptiles and amphibians) since 1970. These substantial population declines compound with large initial populations: mammals were roughly 40 times more numerous than humans in

1970, fish were a million times more numerous, and insects were several billion times more numerous. The disappearance of this large number of sentient animals represents a potentially significant reduction in aggregate sentience.

We find substantial losses for non-human life. We first focus on aggregate sentience, which has the advantage of not requiring us to take a stance on the welfare weight of sentient animals relative to humans. In that case, animal sentience losses outweigh population gains to humans by 4 to 7 orders of magnitude. That is, animal sentience losses are 10,000 to 10^7 times larger than the sentience gains accruing to humans from population growth, or in fact similarly large relative to the sentience of the current human population.

This range represents a suite of variations around our central specification. We vary the parameters scaling the relationship between neurons and cognition around our point estimates. In particular, we also consider piecewise scalings that further downweight all animals relative to humans by leveraging cognitive tests that are specific to primates and humans only and thus may better capture humans' singular ability for cognition and, potentially, sentience. At our most conservative estimate of cognitive scaling within animals, the losses to fish alone are larger than human gains. At our most conservative estimate of cognitive scaling between animals and humans, the losses to animals exceed human gains by multiple orders of magnitude. We also include or exclude invertebrates which have extremely large initial populations and whose moral status is open to disagreement. Yet, even with conservative cognitive scalings and without invertebrates, sentience losses to animals represent 10,000 times the gains to humans.

Including the substantial increase in the populations of farmed animals does not change these conclusions. Farmed animal sentience is of the same order of magnitude as human population sentience. Although farmed animals represent approximately 40% of current earthborne vertebrate biomass, biomass does not scale one-for-one with sentience. The large declines from initially large populations of smaller but still sentient wild animals are quantitatively much more significant. Adjusting farmed animal sentience for their low quality of life would further reduce their contribution to sentience increases.

Converting sentience into welfare requires us to take a stance on the welfare weight of animals relative to humans. The strict utilitarian social welfare function weighs all sentience equally and assigns a welfare weight of 1 to animals, so that aggregate welfare

losses are equal to aggregate sentience losses. We also consider a social welfare function that discounts animal welfare beyond their degree of sentience. When the welfare weight of non-human animal life is zero, other animals do not matter and one recovers a standard welfare function. Varying the welfare weight between 0 and 1 traces out the range of possible valuations of biodiversity loss. When the animal welfare weight is between 10^{-7} and 10^{-4}, the loss of animal life exactly equals the value of human gains depending on the specification.

Quality of life, a species' realized welfare relative to its welfare range, also contributes to aggregate welfare. Our baseline exercise assumes that humans and all other animals had the same quality of life in the state of nature. Incorporating the benefits of economic growth for humans makes little difference to our baseline conclusions because welfare gains from population increases outweigh economic gains at standard values of statistical life (VSL). In addition, assuming lower quality of life for animals compared to humans in the state of nature is isomorphic to assigning a smaller welfare weight to them. Given the large aggregate animal sentience losses that we find, one would need to assume that animals are much closer to indifference between life and death to offset these losses.

We convert these values into a willingness to pay for animal life. We start from a human VSL of USD 10 million. With a welfare weight of one and adjusting for life expectancy, a great ape's life is valued at USD 3.2 million, a crow's life at USD 180,000 and a salmon's life at USD 2,100. These values exceed stated preferences for biodiversity protection by several orders of magnitude (Loomis and White 1996). Next, we consider a welfare weight of 10^{-4} that approximately equalizes animal losses and human gains (in the conservative case excluding invertebrates). In that case, a great ape's life is valued at 320 dollars, a crow's life at 18 dollars and a salmon's life at 21 cents. Animal losses and human gains break even despite these smaller values because of the substantial population declines over the last five decades.

These results reveal a large discrepancy between a sentience-based valuation of animal welfare and the valuation implied by human behavior and willingness to pay. Through the lens of our framework, this suggests that humans assign welfare weights to other animals that are substantially below one. We do not argue for a particular value of the animal welfare weight. Rather, our objective is to highlight the importance of this question given the substantial losses in animal sentience over the past half century. In particular, seemingly small departures from a zero animal welfare weight can lead

to large aggregate welfare costs from biodiversity loss.

Breaking even in welfare space, or even further downweighting animals, does not preclude large aggregate dollar-equivalent losses. We define human gains to include the contribution of population gains and economic growth, which thus exceed annual Gross Domestic Product (GDP) at conventional VSLs. At an animal welfare weight of 10^{-4} that approximately equalizes animal losses excluding insects and human gains, the welfare cost from biodiversity loss represents 112% of global GDP in 2025.

Related Literature. A large literature in economics and ecology has attempted to place a value on biodiversity through the instrumental value of ecosystem services: the direct benefits and services that animal and plant life provide to humanity simply by existing. Costanza et al. (1997) outline seventeen such global services, such as soil generation, pollination, being a source of pharmacological discovery through embodied genetic material, and recreation, which are not priced by the market, concluding that natural services may be worth several multiples of global GDP. The literature has expanded this approach (Simpson et al. 1996; Rausser and Small 2000; Daily et al. 2000; Barbier 2007; Bateman et al. 2013; Giglio et al. 2025) and focused on specific species or environments, where identification of instrumental benefits can be more precisely estimated (Frank and Sudarshan 2024; Taylor and Druckenmiller 2022; Frank 2024).

A second strand of literature studies the value of biodiversity or species as an argument of human utility. Krutilla (1967) provides an early articulation of existence value: individuals may value the continued existence of biological variety even when they do not expect to use or directly experience it. Weitzman (1998) and Metrick and Weitzman (1998) propose to assess biodiversity through its aesthetic or informational value to humanity. Loomis and White (1996) collect direct estimates of stated willingness to pay for species protection and study their correlates. This contingent valuation may be interpreted as the direct value humans place on a species' existence as an additional argument of their utility.

Our work is conceptually distinct from both of these strands. Our approach explores the implications of considering other sentient life as valuable in and of itself as a direct bearer of welfare, in the same way that a human life has intrinsic value regardless of any service to other humans.[1] This approach is grounded in a long tradition in moral philosophy reviewed in Singer (1975) and Birch (2024). Our contribution is quantitative: we combine measures of animal population declines with a cardinal metric of sentience

[1] Adhami et al. (2023) consider a related total utilitarian perspective counting human lives only to re-evaluate the gains from growth.

based on recent advances in neuroscience to operationalize these moral philosophy arguments. The resulting measure of the intrinsic value of biodiversity is, to our knowledge, new to the literature.[2]

Closer to our approach and separate from the valuation literature, work in welfare economics treats non-human animals as direct bearers of welfare (Carlier and Treich 2020; Fleurbaey and Van der Linden 2021). Blackorby and Donaldson (1992) and Espinosa and Treich (2021, 2025) include both human and animal utilities in the social welfare function in the context of animal research and farming for meat. The latter explicitly allow for the possibility that captive animals have negative welfare, in particular due to poor captive living conditions. We assume that, in the state of nature, wild animal lives are, on average, not characterized by net negative welfare. Johansson-Stenman (2018) uses survey data from Sweden to show that there is likely widespread public support for assigning at least some moral weight to the welfare of animals. Budolfson et al. (2024) propose using a total utilitarian framework adjusted for quality of life in order to evaluate policies around meat consumption and biodiversity protection, but do not estimate the welfare score of animals. We complement this literature by quantifying sentience species by species and using it to evaluate the consequences of biodiversity loss.

In contemporaneous work, Eden et al. (2026) develop a related method to ours. They also assess the consequence of changes in animal populations using neuron counts to cardinalize animal welfare. Our approaches differ in two critical dimensions. First, we use observed cognitive tests to anchor how sentience relates to neuron counts, while they posit particular values for this scaling. Second, they assume that many wild animals have negative welfare in their state of nature, implying that wild animal population declines improve aggregate welfare. We view these two approaches as complementary: they reflect how various defensible assumptions can lead to a wide range of conclusions, and highlight the need for further work in this area.

[2]The intrinsic valuation of biodiversity is compatible with the existence valuation to the extent humans choose to assign existence value equal to sentience.

2. Conceptual Framework

2.1 Welfare Function

We define total welfare of animals excluding humans at date t as the sum of the welfare of all animal species living at date t:

$$\mathcal{W}_{At} = \sum_{s \in \mathcal{S}} \lambda_s \times u_s(x_{st}) \times N_{st}, \tag{1}$$

where $\mathcal{S}$ is the set of species excluding humans, λ_s is a weight capturing the relative moral status of species s, and $u_s(x_{st})$ is the instantaneous utility of a representative individual of species s at date t, that we allow to depend on a generic vector x_{st}. N_{st} is the number of individuals of species s living at date t. The welfare of humans $s = 0$ at date t is:

$$\mathcal{W}_{0t} = u_0(x_{0t}) \times N_{0t}, \tag{2}$$

where the moral status of humans is normalized to 1.

Welfare functions of the form in equations (1) and (2) have been proposed in many studies of animal welfare (Budolfson et al. 2024). In this section, we review conceptual definitions of each component of the welfare function in (1) and (2) that are general enough to apply, at least in principle, to many species. We postpone our discussion of the potential comparability between human and animal welfare to Section 2.4. We propose empirical definitions of these components in Section 3.

2.2 Moral Status and Sentience

Only entities with moral status, that is, those with intrinsic moral worth, enter the social welfare function. A broad consensus in moral philosophy, and in particular in utilitarianism, argues that *sentience* is both necessary and sufficient for moral status $\lambda_s > 0$ (Singer 1975, 1979; Hare 1981; DeGrazia 1996; Birch 2024).

The concept of sentience is the natural generalization of the human ability to experience positive or negative feelings and emotions. Formally, sentience is the capacity for valenced experience. In this context, experience means phenomenal consciousness, that is, "there is something it is like for the subject to be in that state." Experience excludes unconscious processing, such as reflexes. Valence is the property of "feeling bad or

feeling good to the subject." Pain and pleasure are two obvious examples of valenced experiences.

This position, known as sentientism, follows naturally from utilitarianism, the moral foundation of modern welfare economics. Indeed, in the consequentialist-utilitarian framework, valence is the relevant scale for moral evaluation. The objective is to maximize the good, where the good is understood as positively valenced experience and the absence of negatively valenced experience across all individuals. Which species these individuals belong to is incidental to utilitarianism. Including non-human sentient beings in the social welfare function amounts to a consistency requirement.[3]

Of course, reasonable individuals can disagree about the implication that species other than humans have moral status. Nonetheless, assigning at least some moral weight to the welfare of non-human living beings strikes us as a widely held view: the existence of animal welfare laws, the widespread care people demonstrate for pets, and the numerous protections for endangered species all point in this direction.

Sentience and cognition. One may reasonably ask whether cognition should also determine moral status. Cognition is not typically regarded as directly relevant for moral status in most contemporary ethical theories, including sentientism.[4] Cognition may nevertheless be indirectly relevant for welfare: it may play a mediating role by amplifying or dampening affective intensity, or expanding the variety and structural richness of valenced experience. We return to the relationship between cognition and sentience in Section 3.

2.3 Utility and Sentience

Having defined who may enter the social welfare function, we define what constitutes welfare. We adopt the mental-state theory of welfare, which posits that welfare consists of consciously experienced valenced states (Crisp 2006; Feldman 2002, 2004). Under this view, the (human) utility function represents specific valenced experiences: a positively valenced experience represents higher utility, and a negatively valenced experience represents lower utility.

The mental-state theory of welfare corresponds to the standard consequentialist-

[3]Sentientism also emerges from other influential moral theories such as Korsgaard's neo-Kantianism and Nussbaum's capabilities approach, despite starting from fundamentally different premises. We discuss classical and preference utilitarianism, as well as these alternative theories, and whether they are consistent or inconsistent with sentientism, in more detail in Appendices A.1 and A.2.

[4]Only certain deontological traditions such as orthodox Kantian ethics tie moral status directly to rational agency. However, this view has been shown to lead to important paradoxes.

utilitarian framework.[5] It is a generalization of classical hedonism, as it relies on all valenced affective states (fear, anxiety, social reward, playfulness, etc.), as opposed to only pain and pleasure.[6]

While it was originally developed in the context of human welfare, the mental-state theory extends naturally to other animals because it relies on sentience rather than specifically human traits, and is thus consistent with sentientism (Singer 1975; DeGrazia 1996; Birch 2024). To the extent that many species exhibit sentience, they also have a utility function u_s that represents the strength of their valenced experiences.

Utility, moral status, and welfare weights. Without further restrictions, degrees of moral status λ_s are isomorphic to degrees of sentience embedded in the utility function u_s. Therefore, our baseline analysis starts by normalizing $\lambda_s = \mathbb{1}\{\sup u_s > 0\} \in \{0, 1\}$ to be an indicator variable encoding whether a living being is sentient at all. Moral status then indicates whether a species has any sentience at all, and the utility function measures how much sentience this species has.[7]

We additionally assess the consequences of assigning sentient animals a welfare weight strictly below 1, further discounting their weight in the social welfare function relative to their measured sentience. To do so, we consider cases in which the utility function u_s is fixed to our cardinal measure of sentience, and animals other than humans have $\lambda_s \in (0, 1)$, where λ_s now plays the role of the relative welfare weight that this species receives relative to humans holding sentience fixed.

Welfare ranges. The utility function defines the capacity for welfare, sometimes called the welfare range, as the total range between its maximum and minimum utility:

$$\overline{u}_s = \sup_{x_{st} \in \mathcal{X}_s} u_s(x_{st}) - \inf_{x_{st} \in \mathcal{X}_s} u_s(x_{st}), \tag{3}$$

where the range spans a set of possible vectors $\mathcal{X}_s$ that might include food quantity and quality, safety, comfort, novel experiences, and so on.[8] We further normalize $\inf_{x_{st} \in \mathcal{X}_s} u_s(x_{st}) = 0$ for all s. This normalization imposes that all species achieve

[5] In particular, the value of an action depends solely on its consequences for individual welfare. That is, there are no other morally relevant properties such as rights, freedom, equality, except insofar as they affect welfare.

[6] The mental-state view is also fully compatible with a preference-satisfaction view of welfare, in its evidential reading: satisfaction of preferences is evidence of welfare, and welfare is valenced experience. We provide a more detailed discussion in Appendix A.1.

[7] This normalization corresponds to the so-called unitarian view. Alternatively, the hierarchical view loads differences in degrees of sentience into differences in moral status as opposed to differences in the utility function. For our purposes, these two views are isomorphic.

[8] We assume that either the utility function is bounded or the set $\mathcal{X}_s$ is compact.

weakly positive utility. We return to the possibility of negative utility in Section 5.5.

We can then rewrite the utility function as $u_s(x_{st}) = \overline{u}_s \times \tilde{u}_s(x_{st})$, where $\tilde{u}_{st} \equiv \tilde{u}_s(x_{st})$ captures the difference between *capacity* for welfare and *realized* welfare. The component $\tilde{u}_{st}$ captures the quality of life with respect to potential.

2.4 Species Welfare and Total Utilitarianism

Having defined sentience, the utility function and moral status, we consider the welfare of a single species, defined as:

$$\mathcal{W}_{st} = N_{st}\, u_s(x_{st}). \tag{4}$$

The formulation in equation (4) continues to build on the standard consequentialist-utilitarian framework, whereby we can rank states of the world by the sum or average of individual welfare. This property requires inter-individual cardinal comparability of welfare and additive aggregation across individuals.

The formulation in equation (4) adopts total rather than average utilitarianism. That is, the social welfare of a species is the sum of the welfare of all individuals of that species, not the average. Total utilitarianism is the natural setup for our investigation. Asking about the welfare effects of the loss of animal populations is only relevant if the welfare criterion is not invariant to the scale of populations. In addition, a total utilitarian social welfare function with variable population emerges from seemingly desirable axioms, as shown in Kuruc et al. (2022).[9]

Aggregating species welfare (4) into total animal welfare (1) involves one additional key property: welfare is comparable across animal species in addition to being comparable across individuals. The mental-state theory and sentientism offer a natural justification for this property: sentience is, in principle, comparable across individuals and species.[10]

[9]The first axiom is 'same number Pareto': the welfare ordering respects Pareto improvements among fixed populations. The second axiom is 'non-anti-egalitarianism': society does not prefer inequality. The third axiom is 'mere addition': holding other individuals' utilities constant, adding a person who values living does not reduce social welfare.

[10]Of course, instantaneous utility differs from lifetime utility. Because humans tend to live longer than most animals, the lifetime utilities of humans and other animals differ more than their welfare capacities. However, our welfare function does not require us to adjust for life expectancy by species because it is defined as flow utility for each living individual.

One may aggregate animal welfare (1) and human welfare (2) into total welfare:

$$\mathcal{W}_t(\boldsymbol{\lambda}) = \mathcal{W}_{0t} + \mathcal{W}_{At}(\boldsymbol{\lambda}), \tag{5}$$

which further requires that welfare is comparable across animals and humans. This property carries the same content as requiring that welfare is comparable across animal species. Of course, comparability of welfare does not imply equal weights on sentience: our formulation allows for smaller welfare weights assigned to animals $\lambda_s < 1$. Changing the vector $\boldsymbol{\lambda} = \{\lambda_s\}_{s \geq 1} \in [0,1]^{|\mathcal{S}|}$ traces out the Pareto frontier defined by the welfare function (5).

To isolate the role of sentience from the role of welfare weights, we consider two comparisons. First, we compare aggregate animal sentience, defined as $\mathcal{W}_{At}(\mathbf{1})$, to aggregate human sentience $\mathcal{W}_{0t}$, holding quality of life fixed. This first comparison has the advantage of abstracting from the welfare weights assigned to animals relative to humans, about which one may reasonably disagree. Of course, it only compares units of sentience and does not translate into an actionable welfare metric. Second, we map out how aggregate animal welfare $\mathcal{W}_{At}(\boldsymbol{\lambda})$ compares to aggregate human welfare $\mathcal{W}_{0t}$ as a function of animal welfare weights $\boldsymbol{\lambda}$ and economic growth embedded in $\tilde{u}_{0t}$. This second comparison requires us to take a stance on animal welfare weights $\boldsymbol{\lambda}$.

Having established the conditions under which the welfare function (5) emerges, we now turn to the measurement of each of its components.

3. Measuring Sentience

3.1 Measurement Challenges

Our exercise requires the measurement of two objects for a large number of species: moral status λ_s, strictly positive if species s is sentient, and $\overline{u}_s$, a cardinal measure of the range of valenced experiences of species s. These objects are inherently difficult to measure: conscious experience gives subjects a first-person point of view that cannot be fully apprehended from outside but only from the inside. This property constitutes the 'hard problem' of consciousness (Chalmers 1995).

Our measurement approach relies heavily on the established link between sentience and the brain. The overwhelming majority of contemporary philosophers and neuroscientists hold the view that consciousness and valenced experience are ultimately

physical phenomena, realized in physical processes occurring in the brain. Mental states, such as thoughts, feelings and subjective experiences, arise from neural activity, and a complete physical account of the brain is also a complete account of mental states (Papineau 2002; Kim 2005).[11] Therefore, variation in neural architecture is a principled starting point for understanding interspecies variation in sentience and welfare ranges.

Empirical frameworks for measuring sentience have emerged as the question of animal sentience has migrated from scientific and bioethical inquiry toward explicit legal recognition over the past three decades. The first systematic criteria emerged from the 1991 report of the Institute of Medical Ethics working party on the use of animals in biomedical research (Smith and Boyd 1991). The European Union committed to regard animals as "sentient beings" in Article 13 of the Lisbon Treaty in force since 2009, building on an earlier protocol to the 1997 Treaty of Amsterdam. The United Kingdom enshrined sentience in domestic law through the Animal Welfare (Sentience) Act 2022. The legal landscape is actively evolving, and we attempt to build on the most up-to-date conceptual framework. In particular, we borrow heavily from the framework in Birch (2024), used as the basis of that Act.

3.2 Behavioral Markers of Sentience

The literature defines markers of sentience. These markers, while individually neither necessary nor sufficient to determine sentience, are to be understood as "symptoms" of sentience, each raising the probability that an individual is sentient. They are useful not only to determine sentience, but also to identify which brain regions are responsible for valenced experience.

The starting point is humans, who can report experiencing a valenced experience at the same time as displaying a particular marker. These markers are then generalized to other animals by analogy. The literature has converged on three types of markers: behavioral markers such as avoiding a particular location, physiological markers such as an autonomic stress response, and pharmacological markers such as reaction to an analgesic. Following Birch (2024), these criteria are:

1. **Motivational trade-offs:** The animal engages in dynamic decision-making, weigh-

[11] This view is called materialism. The main alternative is dualism, which posits that to explain consciousness, we need to posit at least some fundamentally non-physical properties. There is, however, no scientific basis for this claim. It is beyond the scope of this paper to settle this debate, and we accept the materialist position. Appendix B.1 provides a brief account of the debate and of the arguments that have made materialism the contemporary mainstream.

ing the adverse impacts of noxious stimuli against the value of rewards.

2. **Flexible self-protection:** The animal exhibits flexible self-protective behaviors, including wound tending and guarding.
3. **Associative learning:** The animal demonstrates associative learning by forming connections between noxious stimuli and neutral cues.
4. **Analgesia preference:** The animal expresses a preference for analgesics or anaesthetics when injured.

These markers are useful to define the scope of sentience. However, comparable and cardinal measures of these markers are not available for a broad range of species. Thus, for the purpose of measuring welfare ranges in a systematic way across species, we turn to the brain regions that are related to behavioral markers of sentience.

3.3 Neurological Markers of Sentience

The neural basis of sentience. Two broad categories of brain structure have been identified to be most likely responsible for valenced experience and thus the markers listed above. The first are evolutionarily ancient subcortical structures, in particular the periaqueductal gray and the mesolimbic dopamine system. We will refer to these structures as the 'midbrain' for brevity. The second are evolutionarily more recent integrative brain regions: the cortex in mammals, and functional analogs in other taxa. We will refer to these structures as the 'pallium' for brevity.

The midbrain is generally understood to govern anoetic consciousness: the most basic notion of phenomenal consciousness that does not involve self-referential content. The pallium contributes to noetic and autonoetic consciousness, which are more complex forms of phenomenal consciousness. They respectively involve the recognition that one is experiencing and the embedding of that experience in time. While specific theories disagree about whether anoetic, noetic or autonoetic consciousness are necessary to achieve sentience, all sides agree that sentience emerges from the architecture of the brain.[12]

[12]Some theories hold anoetic consciousness to be sufficient for valenced experience, with pallium mechanisms further elaborating this experience (Panksepp 1998; Merker 2007; Damasio and Carvalho 2013). Competing theories hold that subcortical activity is unconscious and needs cortical re-representation to become an experience of any kind. These theories include, among others, the Global Neuronal Workspace theory (Dehaene et al. 2017), higher-order theory (LeDoux and Brown 2017; Rosenthal 2005), and recurrent processing theory (Lamme 2006). All locate consciousness in different brain mechanisms but share the requirement of cortical involvement. We provide more details in Appendix B.2.

For the determination of sentience based on brain architecture, Birch (2024) proposes four neurological markers. As for the behavioral markers discussed in Section 3.2, none of these markers is strictly necessary or sufficient for sentience. Rather, each criterion that is satisfied raises the probability that the animal is sentient. These criteria are:

5. **Nociceptors:** The animal possesses receptors located in neurons that respond specifically to noxious stimuli.
6. **Integrative brain regions:** The animal possesses brain regions capable of integrating information from various sensory sources.
7. **Integrated nociception:** Neural pathways within the animal link nociceptors to integrative brain regions.
8. **Analgesia:** Behavioral responses to noxious stimuli are modulated by chemical compounds affecting the nervous system.

The evolutionary basis of sentience. The neural architecture responsible for sentience is the product of evolution. These evolutionary roots motivate why sentience is likely to be shared by many species. The leading theories of sentience treat valence as an evolved adaptation that serves to motivate fitness-improving actions and discourage fitness-reducing ones (Ginsburg and Jablonka 2019; Godfrey-Smith 2017). Under this view, valenced experience plays three roles: it represents fitness-relevant information, it serves as a common 'currency' in the brain for decision-making, and it facilitates learning by tagging outcomes as good or bad. Therefore, animals with shared evolutionary history are likely to have similar valenced responses, with similar scopes for welfare experience supported by similar nervous systems (Browning 2019).

3.4 Sentience Candidates

Because sentience can rarely be established with complete certainty, we adopt Birch (2024)'s definition of a sentience candidate, and a candidate for moral status: "A system S is a sentience candidate if there is an evidence base that (a) implies a realistic possibility of sentience in S that it would be irresponsible to ignore when making policy decisions that will affect S, and (b) is rich enough to allow the identification of welfare risks and the design and assessment of precautions."

Applying the markers 1–8 to the available evidence, Birch (2024) classifies as sentience candidates all vertebrates (mammals, birds, fish, reptiles, and amphibians), adult coleoid cephalopods in the phylum Mollusca, decapods in the suborder Pleocyemata

of the phylum Arthropoda, and adult insects in the phylum Arthropoda, though he notes that the evidence is sparser for insects.

By contrast, decapods of the Dendrobranchiata suborder, including penaeid shrimps, together with insect larvae, spiders, gastropods, and nematode worms are instead "investigation priorities": they fall short of sentience candidature, but further evidence could plausibly elevate them to it. Plants and unicellular organisms are not sentient.

3.5 Welfare Ranges and Cognition

With an empirical definition of sentience linked to brain architecture, we are now ready to construct a cardinal measure of sentience for each species. While the moral status of different species has received considerable attention, the existing literature does not provide the cardinalization of welfare ranges that our exercise requires.

To construct a cardinal measure of welfare ranges $\overline{u}_s$ for a large number of species, we start from the idea that species with similar brains will have similar welfare ranges, and hence variation in neural architecture is a principled starting point for understanding interspecies variation in welfare ranges. Following Section 3.3, we take $\overline{u}_s$ to be a function of brain characteristics, which are observed for many species. We posit the general relationship that allows for the midbrain and the pallium to be relevant for sentience, and thus, welfare ranges:

$$\overline{u}_s = (z_{m,s}^{\alpha_u} z_{p,s}^{1-\alpha_u})^{\kappa_u}, \tag{6}$$

where $z_{m,s}$ is the number of midbrain neurons of species s and $z_{p,s}$ is the number of pallial (cortical) neurons of species s. α_u captures the relative contributions of midbrain and pallial neurons to welfare ranges. κ_u captures the overall scaling in neural architecture.

The problem then reduces to estimating the elasticity of welfare range with respect to brain characteristics. The central question for interspecies comparison is the curvature of this relationship: Do welfare ranges exhibit decreasing, constant, or increasing returns to brain characteristics?

Measuring $\overline{u}_s$ directly to fit this function is not feasible. Existing tests for sentience typically use experimental protocols that are not designed to be comparable across species. Sentience indices typically aggregate binary tests by relying on "scorecard approaches" that take values in a bounded range by construction, which would artificially

compress welfare ranges.

We circumvent this issue by relying on measures of cognition across species. Standardized cognitive tests yield continuous performance measures that are comparable across species. We can thus fit the relationship between test performance and brain characteristics for the species for which such tests exist, and extrapolate it to the remaining species. Cognition is not sentience, as detailed in Section 3.3. However, we claim that using cognition is conservative: the neuroscientific evidence suggests that there is wider dispersion in cognition than in sentience across species, and hence relying on cognition is more likely to understate than to overstate the welfare range of non-human animals. That is, measures of observed cognition will allow us to provide a lower bound on the welfare ranges of non-human animals relative to that of humans.

To formalize the content of this claim, we write cognition as a function of midbrain and pallial neurons, similarly to welfare ranges:

$$c_s = (z_{m,s}^{\alpha_c} z_{p,s}^{1-\alpha_c})^{\kappa_c}, \tag{7}$$

where the parameters α_c, κ_c may differ from their counterparts for welfare ranges.

Total neuron count is $z_{tot,s} = z_{m,s} + z_{p,s} + z_{o,s}$, where $z_{o,s}$ is other neurons (e.g. the cerebellum). For expositional simplicity, and as we show later empirically in Figure C.4, we posit that neuron counts in each brain region are log-linear in total neuron count:

$$z_{p,s} = z_{tot,s}^{\gamma_p}, \qquad z_{m,s} = z_{tot,s}^{\gamma_m}.$$

3.6 Cognition as a Lower Bound on Welfare Ranges

We are now ready to state the main assumption that lets us construct a cardinal lower bound on sentience.

Assumption 1 (Sentience and Cognition)**.**

1. *Pallial neurons grow faster with total neuron counts than midbrain neurons:* $\gamma_p \geq \gamma_m$.
2. *Cognition depends relatively more on the pallium than sentience:* $\alpha_u \geq \alpha_c$.
3. *Cognition scales more strongly with neuron counts than sentience:* $\kappa_u \leq \kappa_c$.

Below, we present the main arguments supporting these assumptions in a concise and nontechnical form. A more detailed discussion of the scientific evidence underlying these assumptions is provided in Appendix B.4.

The first part of Assumption 1 is that the pallium grows relatively faster with brain size than the midbrain. This assumption reflects the empirical evidence (e.g., Kverková et al. 2022). Indeed, through evolution, increases in brain size are accompanied by disproportionate expansion of the pallium, while the midbrain remains comparatively conserved, causing pallial neuron counts to scale more steeply with total neuron counts.

The second part of Assumption 1 requires that sentience weighs the midbrain more heavily than the pallium, relative to cognition. Neuroscientific evidence indicates that cognition (working memory, abstract reasoning, and planning) occurs primarily in the cortex. By contrast, sentience relies heavily on the midbrain for all anoetic, noetic and autonoetic theories of sentience (Panksepp 1998; LeDoux and Brown 2017). In addition, sentience is thought to predate flexible cognition evolutionarily, likely implying that it relies more on older brain structures (Panksepp 1998; Damasio 1999; Damasio and Carvalho 2013).

The third part of Assumption 1 encodes the idea that cognition depends weakly more on overall neuron counts than sentience. Neuroscientific evidence suggests that cognition is strongly linked to total neuronal capacity (Herculano-Houzel 2017; Dicke and Roth 2016). Cognitive activity is associated with high-dimensional neurological signals, whose complexity rises with neuron count without obvious anatomical limits (Rigotti et al. 2013; Stringer et al. 2019). By contrast, valenced experience is associated with low-dimensional neurological signals (Levy and Glimcher 2012). These signals are in turn associated with neurochemical reactions that eventually saturate (Koob and Volkow 2010). These properties are consistent with the role of valence as a common scale for decisions.

With Assumption 1 in hand, we state our main result.

Proposition 1. *(Cognition as a Lower Bound on Sentience) Suppose that humans $s = 0$ have the largest cognition. Under Assumption 1, for all species $s \in \mathcal{S}$:*

$$\frac{d \log \overline{u}_s}{d \log z_{tot,s}} \leq \frac{d \log c_s}{d \log z_{tot,s}} \qquad \textit{and so} \qquad \frac{c_s}{c_0} \leq \frac{\overline{u}_s}{\overline{u}_0}.$$

Proof. $\frac{d \log \overline{u}_s}{d \log z_{tot,s}} = \kappa_u \left(\alpha_u \gamma_m + (1-\alpha_u)\gamma_p\right) \leq \kappa_c \left(\alpha_u \gamma_m + (1-\alpha_u)\gamma_p\right) \leq \kappa_c \left(\alpha_c \gamma_m + (1-\alpha_c)\gamma_p\right) = \frac{d \log c_s}{d \log z_{tot,s}}$, where the first inequality uses Assumption 1.3 and the second inequality uses Assumption 1.1 and 1.2. The second part of the Proposition follows from integrating the first part between $z_{tot,s}$ and $z_{tot,0}$. □

Proposition 1 formalizes the idea that cognition is a lower bound for sentience. It is

conservative in that several of the inequalities involved in the proof of Proposition 1 are likely to be strict.

Of course, Proposition 1 is useful only if we have access to a cardinal measure of cognition. If we observe cognition for all species, there is no further requirement. If we observe cognition for a subset of species, we also need to observe neuron counts for this same subset in order to extrapolate to all species. We thus require:

Assumption 2. *(Cognition is Observable) Either one of the following conditions holds:*

1. *Relative cognition* $\frac{c_s}{c_0}$ *is observable for all species* $s \in \mathcal{S}$.
2. *Relative cognition* $\frac{c_s}{c_0}$ *and total neuron counts* $z_{tot,s}$ *are observable for a subset of species* $s \in \overline{\mathcal{S}} \subseteq \mathcal{S}$.

Assumption 2 emphasizes that choosing a measure of cognition sets the scale of the lower bound. Given any measure of cognition, a transformation thereof leads to different values of the lower bound. Thus, choosing a measure of cognition is an important choice and is not innocuous. For this reason, we explore how our results change when we vary the normalization of cognition metrics.

In practice, we observe cognition measures only for a subset of species. We rely on Assumption 2.2 and leverage the relationship between cognition and total neuron counts. Therefore, we estimate relationships of the form:[13]

$$\log c_s = \log a + \beta \log z_{tot,s} + \epsilon_s \quad , \quad s \in \overline{\mathcal{S}} \quad , \tag{8}$$

and use the estimated coefficient β to extrapolate to all species. Combining an estimate of β from equation (8) with Proposition 1, we obtain our lower bound:

$$\frac{\overline{u}_s}{\overline{u}_0} \geq \left(\frac{z_{tot,s}}{z_{tot,0}} \right)^{\beta} \quad , \quad s \in \mathcal{S}. \tag{9}$$

Equipped with an empirical bound on welfare ranges, Section 4 discusses the animal population data, the cognition data, and the neuron count data that we use to estimate

[13] When the relationships of midbrain and pallial neurons to total neurons have perfect predictive power with exactly the same exponents across all species, we can equivalently consider the elasticity of cognition with respect to total or pallial neurons and use the predicted values. In practice, the R-squared values of these relationships are not exactly one, and the scaling exponents do vary slightly across taxa. In particular, γ_p is somewhat larger for primates than for other mammals due to a larger number of pallial neurons relative to total neurons for this order. In that case, it may be more conservative to estimate the elasticity of cognition with respect to pallial neurons, which are the best predictor of cognitive performance across species (MacLean et al. 2014; Herculano-Houzel 2017). We use both strategies in practice and obtain similar results.

all the components of the welfare cost of biodiversity loss.

4. Data and Measures of Sentience

4.1 Population Data

We begin by collecting and harmonizing data on population-level trends for wild animals at different levels of aggregation. We restrict our analysis to vertebrates (Chordata), molluscs, and arthropods, as these correspond to the criteria for sentience candidates and moral status as discussed in Section 3.4. We construct a hierarchical population dataset at the species, order, and group level. We combine estimates of current population levels with estimates of relative population change over time to estimate the absolute change in population size.

Current population levels. For mammals and birds, we obtain species-level population counts from Greenspoon et al. (2025) and Wikipedia's compilations of species counts after verifying its sources. We obtain group-level population totals for birds, fish (mesopelagic and non-mesopelagic), and arthropods from Bar-On et al. (2018). We obtain a total count for mammals by aggregating species-level counts. We obtain a separate count for insects from the Smithsonian Institution (1996). For herptiles, we rely on the estimates in Tomasik (2018). We provide validations of these estimates in Appendix C.1.3.

Cephalopods (belonging to molluscs) and decapod crustaceans (belonging to arthropods) belong to our list of sentience candidates. However, no source allows us to enumerate their populations within the broader arthropod and mollusc taxa. Therefore, they are dropped from our analysis.

Population changes. Our main data source is the data underlying the Living Planet Index (LPI). This data contains time-series population data for 29,559 populations across 4,844 vertebrate species, from 1970 to the present. We replicate the LPI methodology to aggregate these population time series for each major vertebrate group (mammals, birds, fish, herptiles), as well as at the order level for well-covered orders (supported by more than 100 population series). The LPI has a number of known caveats, discussed in Appendix C.1.1. To validate our estimates of population changes at the group and order level, we conduct extensive cross-checks of the LPI-estimated changes with other sources (see Table C.3). This analysis shows that LPI-estimated changes have the same

TABLE 1: Current Populations of Major Animal Taxa

TAXON	CURRENT POP.	HISTORICAL POP.	DECLINE (%)
Birds	3.00×10^{11}	$5.49\ (3.78–7.76) \times 10^{11}$	45 (21–61)
Fish	1.51×10^{15}	$2.84\ (2.26–3.70) \times 10^{15}$	47 (33–59)
Herptiles	1.00×10^{14}	$9.96\ (3.62–28.9) \times 10^{14}$	90 (72–97)
Insects	1.00×10^{19}	1.82×10^{19}	45
Mammals	7.84×10^{10}	$14.5\ (8.7–26.3) \times 10^{10}$	46 (10–70)

Notes: Current population refers to 2020. Historical population refers to 1970. Confidence bounds for the LPI estimates use the bootstrap procedure defined in the LPI technical documentation.

order of magnitude as alternative sources.

The LPI does not cover invertebrates. We obtain an estimate of the insect population change from Dirzo et al. (2014). We complement the LPI with species-level population changes for some mammals, obtained from Greenspoon et al. (2025). Finally, for farmed animals, we use data from the Food and Agriculture Organization Corporate Statistical Database (FAOSTAT).

Final population dataset. We combine present population counts and changes to estimate animal population losses since 1970. We perform this calculation in two ways. Our first approach prioritizes the most granular level of disaggregation: we use species-level counts and changes when available for mammals, otherwise fall back to order-level changes, which are valid only for mammals and birds, and then fall back to changes for the major taxon when a species is not covered at the species or order level. Our second approach uses only changes and counts for the major taxa. Reassuringly, both approaches deliver similar results.[14]

Table 1 summarizes the data at the level of the five major taxa we consider. The most numerous group is insects, outnumbering the other groups by many orders of magnitude. Mammals are the least numerous group, on the order of 10^{10}. These substantial disparities in baseline levels are an important quantitative determinant of the results below. Therefore, we explore the robustness of our conclusions to excluding insects for our welfare calculations.

Table 1 also reports population changes at the taxon level since 1970.[15] All taxa, from the least to the most numerous, have experienced large population declines. Figure 1

[14]For fish, the LPI population series mostly cover fish species living less than 200 meters deep. There is very little data on mesopelagic fish, living deeper in the sea. To be conservative, we assume that the LPI population decline applies only to non-mesopelagic fish, and assume no change for mesopelagic fish. In robustness checks, we present results where we assume the LPI population change applies to all fish.

[15]Figure C.2 depicts these declines graphically.

FIGURE 1: Population Declines for Vertebrates: Order Level

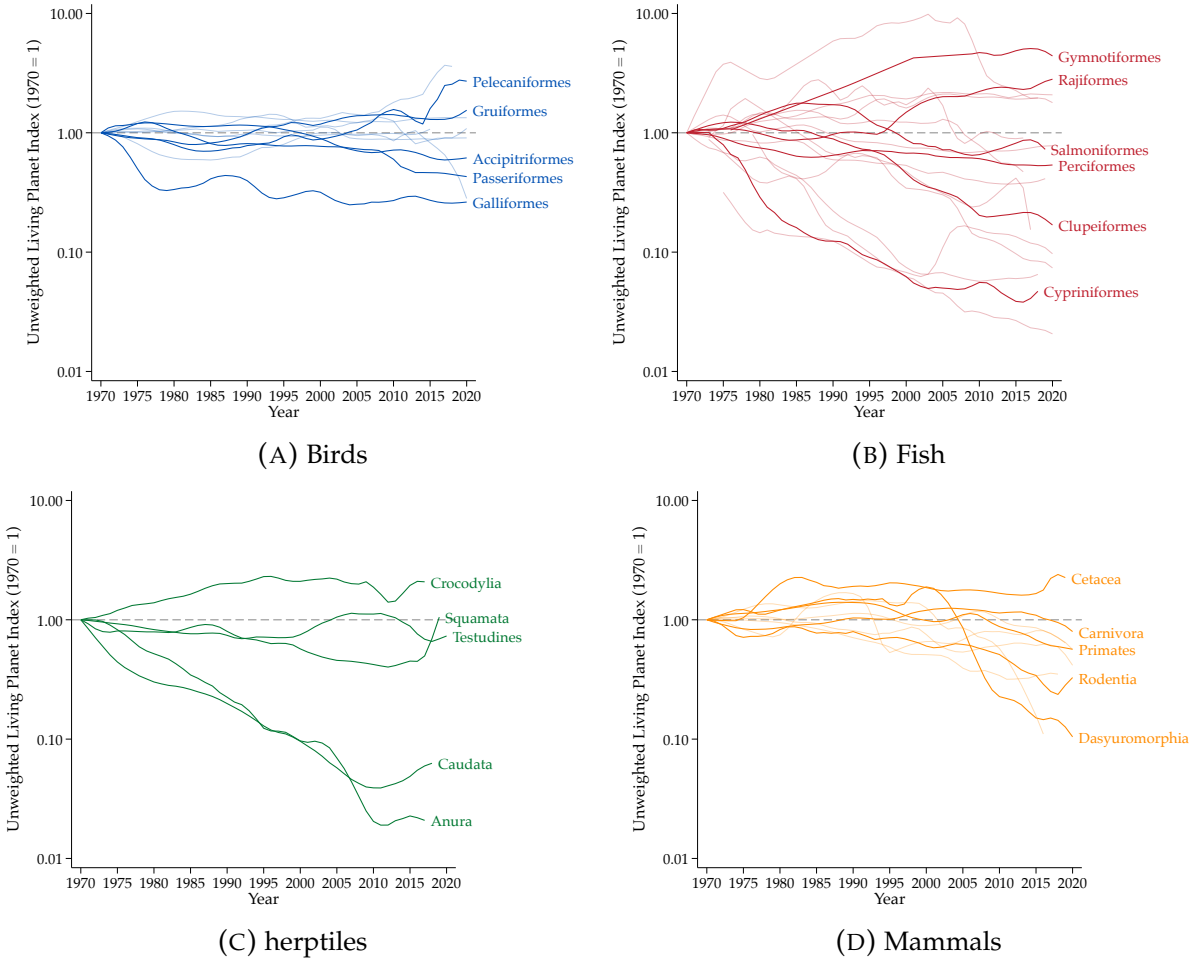

(A) Birds (B) Fish

(C) herptiles (D) Mammals

Notes: Global order-level population trends for Birds, Fish, herptiles and Mammals. Data restricted to orders with more than 100 population time series.

plots population declines at finer taxonomic levels. Most orders have seen substantial population declines since 1970, ranging from 5 to 10% for carnivores and primates to 90% for frogs and salamanders. Cetaceans and Crocodylia are among the few orders to have increased in population, the former likely due to the cessation of commercial whaling from an initially low population in 1970.

These losses were most concentrated in particular areas on the planet. In order to spatially disaggregate the data, we use range maps from the International Union for Conservation of Nature (IUCN) at the species level, and assume population changes are uniform across the species range. For our major groupings, Figure C.3 plots the spatial concentration of decline. Losses of arthropods and mammals were concentrated in Sub-Saharan Africa and South America, while Southeast Asia saw the biggest declines of birds. These are regions of the planet that since 1970 have recorded substantial poverty reductions and economic growth (particularly in Southeast Asia, with more

recent gains since 2000 in Sub-Saharan Africa and South America).

By contrast, farmed animals have seen population increases since 1970, particularly for chickens (Figure C.1). Moreover, this increase is widespread across the different regions of the globe.

Despite our efforts to validate the estimates of population levels and decline across taxa (see Appendix C.1), we want to emphasize that there remains considerable uncertainty about both of these variables. This reflects both genuine scientific disagreement and the difficulty of measuring populations for species that are not a conservation or management priority of humanity. We take the best available evidence, but as the science progresses, the exact magnitudes resulting from applying our methodology will change by necessity.

4.2 Neuron Counts

With measures of animal populations in hand, we collect data on both the total number of neurons and the number of pallial neurons across animal species. We harmonize data from a variety of recent studies (Kverková et al. 2022; Herculano-Houzel et al. 2015; Sol et al. 2022; Olkowicz et al. 2016; Marhounová et al. 2019; Ahrens et al. 2013; Teles et al. 2025; Barron and Mourmourakis 2023) and from Wikipedia after verifying its sources. Table C.5 describes these data sources in more detail.

While extensive, these data do not have complete coverage. Table C.6 lists all assumptions we use to interpolate missing data. We detail the main assumptions in this section. Total neurons are more widely available than pallial neurons. Therefore, we rely primarily on total neurons for our analysis. We do not lose much information in doing so, as Figure C.4 indicates that pallial and total neurons follow a tight log-log relationship across a wide range of species with an R-squared of 0.87. In addition, we observe neither total nor pallial neurons for some species. In this case, we appeal to cellular scaling laws. The biological literature finds that within taxonomic orders (but not across), the number of cells and neurons across species is well explained by a power function of body mass (Herculano-Houzel 2011; Herculano-Houzel et al. 2006, 2007; Sarko et al. 2009). Therefore, within each order for which we have at least three species with observed total neurons, we fit a log-log regression of total neurons on body mass. We then use this regression and each species' body mass to predict the total number of neurons for species for which we observe neither total nor pallial neurons.[16]

[16]We take body masses from PanTHERIA (Jones et al. 2009), EltonTraits (Wilman et al. 2014), and

4.3 Cognitive Tests

With measures of neuron counts in hand, we collect information on cognitive performance. We consider three sets of cognitive tests. Our first set of cognitive tests consists of the 'A-not-B task' and the 'cylinder task' (MacLean et al. 2014; Barron and Mourmourakis 2023). These tests have been performed on species from a broad range of taxa, including great apes, lemurs, monkeys, rodents, corvids, other birds, and fish.

The A-not-B task consists of first familiarizing an animal with finding food in one location (location A). In the test trial, the animal sees the food hidden in the same location (location A) but then moved to a new location (location B) before they are allowed to search. The cylinder task consists of familiarizing an animal with finding food hidden inside an opaque cylinder. The cylinder is then replaced by a transparent cylinder. To successfully retrieve the food, the animal needs to inhibit the impulse to reach for the food directly (bumping into the cylinder) and instead take the detour to the open ends. Both the A-not-B and cylinder tasks are cognitively demanding because they require subjects to inhibit a previously reinforced or prepotent response and select an alternative action. Performance depends strongly on the development of cortical, and especially prefrontal, systems supporting working memory and inhibitory control.

The second set of cognitive tests consists of numeration tests. They assess whether animals can discriminate between different quantities. Typically, an animal must distinguish two sets of objects differing in number, for example choosing between 1 and 3 objects. One can increase the difficulty by making the ratio of the two sets of objects closer to 1 (for example 7 vs. 8). These experiments have been conducted on a range of species, ranging from honeybees to chimpanzees. We aggregate 26 different studies and detail the data we use from each one in Table C.9.

The third set of cognitive tests is the Primate Cognition Test Battery (PCTB). It consists of a collection of cognition tests introduced to compare primates and young humans, and recently extended to parrots (Krasheninnikova et al. 2019). The tests are designed to assess a broad range of abilities such as understanding of space, quantities, causality, social learning, and communication (detailed in Table C.8). They assess a broader range of cognitive abilities, but on a narrower range of species.

An important concern is that these tests may not accurately represent cognition because adult humans succeed trivially at them. As a result, projecting success rates on neuron counts would underestimate the gap in cognition between humans and other

FishBase (Froese and Pauly, eds 2026).

species, and thus, sentience.

The case of young humans helps assess this concern. Young humans perform similarly to great apes on the battery of cognitive tests, which suggests that the two have similar cognition. In turn, we typically treat young humans as having similar sentience to adult humans. Nevertheless, we assess the robustness of our results with a piecewise approach that allows for more convexity in welfare ranges at the top of the distribution.

4.4 Neuronal Capacity and Cognition

Equipped with measures of test performance and neuron counts for a subset of species, we are ready to leverage equation (8). We estimate the relationship between cognition and total neuron counts separately for each test, before aggregating. For our baseline, we only consider the A-not-B, cylinder, and numeration tests, which cover a wide array of species and are thus most amenable to estimating the cognition-neuron relationship from insects to humans. We return to the PCTB in Section 5.3. We consider the success rate $c_s \in [0, 1]$ as our baseline measure of cognition. We emphasize that this choice is not scale-free, and that transformations thereof affect our results. We explore robustness to rescaling our cognition measure below.

Figure 2 plots the relationship between log neuron counts and log performance on the A-not-B and cylinder tests at the taxonomic group level. These plots show that cognitive performance is increasing in the number of total neurons. Relating the log success rate at the A-not-B test to the log total neurons as per equation (8), we obtain $\hat{\beta}^{\text{A-B}} = 0.37$, which is statistically different from 0 at the 5% level. The same regression for the Cylinder test leads to $\hat{\beta}^{\text{Cyl}} = 0.09$, although the point estimate is not statistically significant.

At the group level, the relationship appears to be relatively tight. For the A-not-B test, the R-squared is 0.78. Of course, this group-level relationship masks within-group heterogeneity. Figures D.1 and D.2 display the scatter plots at the species level and for pallial neurons. As we weight our group-level regressions by species and group population, the point estimates are comparable in the species-level analysis. If anything, they are lower for the A-not-B test, which turns out to be less conservative than the group-level estimate.

Figure D.3 plots the relationship between total neuron counts and performance on the numeration tests. The estimates of the slope are always smaller than for the

FIGURE 2: Cylinder and A-not-B Score as a Function of Total Neurons

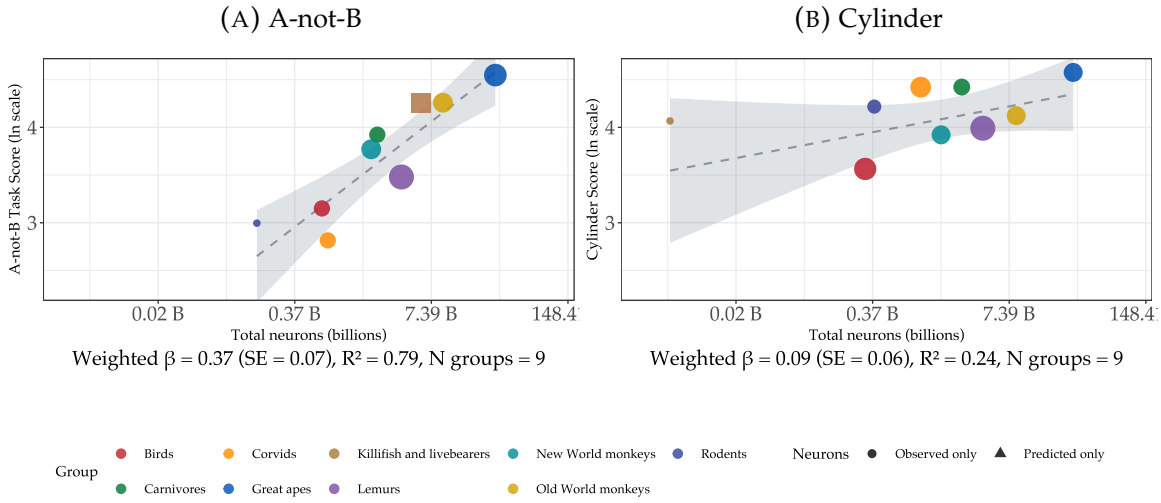

Notes: Panel a: A-not-B test score versus total neuron count (log–log). Panel b: cylinder test score versus total neuron counts (log–log). See Figures D.1 and D.2 for species-level and pallial neurons.

A-not-B task. One reason is that insects, which have very few neurons, tend to perform very well on numeration tests. A larger slope implies a stronger decay of cognition with neuron count. Therefore, we use $\hat{\beta}^{\text{A-B}}$ as a conservative value for the relationship between cognition and neuron counts across species. In what follows, we denote it by β for simplicity.

4.5 Conversion into Utility

Equipped with a baseline estimate of β, we now leverage equation (6) and construct our lower bound on sentience, or welfare ranges, for all groups relative to humans: $\omega_s \equiv \overline{u}_s/\overline{u}_0 = (z_{tot,s}/z_{tot,0})^{\beta}$. Figure 3 displays a histogram of ω_s. We first display this histogram for our baseline estimate $\beta = 0.37$. In that case, many animals have sentience above 0.1 times human sentience.

We also include larger values of β to illustrate its importance. Changing β is isomorphic to rescaling cognition, and thus by varying β we explore the range of alternative cardinalizations of cognition. A smaller value of β implies that all species are similar to humans. A larger value of β implies that species with fewer neurons are substantially less sentient than humans.[17] Varying β equivalently spans the range of values implied by our various cognitive tests. We display the histogram using $\beta = 1$

[17] A few species have more neurons than humans because of body mass extrapolation: the killer whale and long-finned pilot whale. We cap relative welfare ranges of these species at the maximum among species with lower relative welfare range than humans to ensure that they do not drive our results.

FIGURE 3: Relative Welfare Ranges

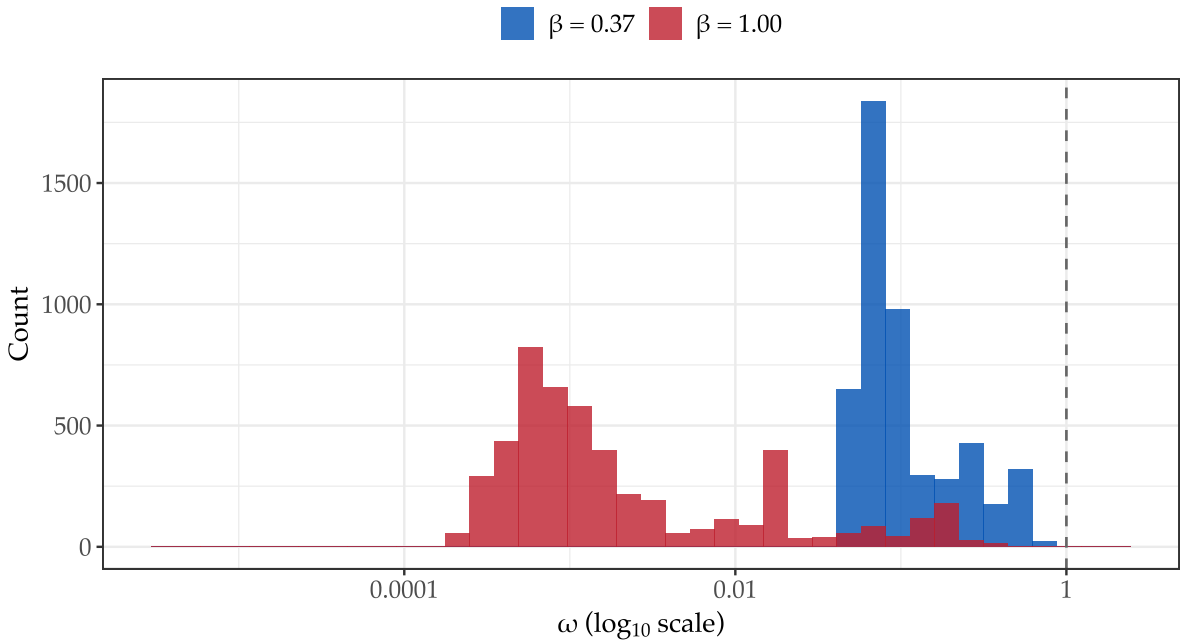

Notes: Histograms of welfare capacity relative to humans $\omega_s = (z_{tot,s}/z_{tot,0})^{\beta}$ for $\beta \in \{0.37, 1\}$. These histograms use observed or predicted neurons at the species level for mammals, means at the order level for birds, and means at the group level for other groups. The dashed line corresponds to $\omega = 1$: welfare capacity equal to that of humans.

as a natural alternative in which sentience is linear in neuron counts. When we use this larger value, the distribution of sentience shifts to the left: few species achieve sentience above 0.1 times that of humans. Most have sentience less than 0.01 times that of humans.

Quality of life adjustment. Each species enters the welfare function with weight $\lambda_s \times \tilde{u}_{st} \times \overline{u}_s$, where the term $\tilde{u}_{st}$ captures realized welfare relative to potential: quality of life. We assume that in a state of nature, denoted $t = 0$, humans and all animals had equal realized welfare relative to potential, i.e., $\tilde{u}_{s0} = \tilde{u}_{00}$ for all species $s \in \mathcal{S}$. Since then, human quality of life has improved dramatically due to economic growth.

We thus proceed in two steps. First, we perform our exercise without accounting for economic growth. In this case, $\tilde{u}_{st}$ is constant for all species s and time period t, including humans. The chosen level of quality of life $\tilde{u}_{s0}$ is inconsequential, since only animal welfare relative to human welfare matters, which turns out to coincide with animal sentience relative to human sentience. In particular, $\tilde{u}_{s0}$ may be arbitrarily small if life in the state of nature is so difficult that wild lives have close to zero net welfare.[18]

[18]There is an active debate over whether the average welfare of wild lives is even positive (Groff and Ng 2019; Browning and Veit 2023). We take positive but arbitrarily low realized welfare, together with symmetry between humans and other animals in the state of nature, as a natural benchmark, and

Second, in Section 5.5, we account for the diverging quality of life trends between humans and animals due to economic growth.

4.6 Merged Population and Sentience Data

The final merged dataset pairs each animal group's population change from 1970 to today with a neuron count, yielding a proxy for its sentience as described above.

Each group is described at a different resolution, matched to how much data is available. Mammals contain one entry per species. Birds contain one entry per order. For both groups, the data construction falls back to order- or class-level averages when more disaggregated data is not available. Fish, herptiles, and Insects contain a single entry at the class level, pairing the aggregate population loss and average neuron counts for these groups. The merged data construction is detailed in Appendix C.3.

5. The Cost of Biodiversity Loss

5.1 Aggregate Sentience Changes

For each group g, we consider the sentience gain or loss between 1970 and t relative to a counterfactual where species s' population remains constant. As described in Section 4.5, the aggregate sentience gain or loss of group g is:

$$\Delta \mathcal{W}_{gt}(\mathbf{1}) = \sum_{s \in g} \overline{u}_s \tilde{u}_{st} \left(N_{s,t} - N_{s,1970} \right).$$

It corresponds to the contribution of group g to welfare if its welfare weight were equal to 1. We then express aggregate sentience losses relative to the gain in aggregate human sentience:

$$\frac{-\Delta \mathcal{W}_{gt}(\mathbf{1})}{\Delta \overline{\mathcal{W}}_{0t}} = \sum_{s \in g} \frac{\overline{u}_s}{\overline{u}_0} \frac{\tilde{u}_{st}}{\tilde{u}_{0t}} \frac{N_{s,1970} - N_{s,t}}{N_{0,t} - N_{0,1970}} \geq \sum_{s \in g} \omega_s \frac{N_{s,1970} - N_{s,t}}{N_{0,t} - N_{0,1970}}, \tag{10}$$

where $\overline{\mathcal{W}}_{0t}$ denotes human welfare holding quality of life fixed at $\tilde{u}_{0t} = \tilde{u}_{00}$, where we assume fixed and equal animal and human quality of life $\tilde{u}_{st} = \tilde{u}_{00}$ for all species s and time periods t. The inequality exploits our lower bound and applies to groups whose populations decline. We introduce the welfare effect of improvements in human quality

examine the sensitivity of our results to this assumption in Section 5.4.

FIGURE 4: Sentience Losses or Gains by Taxa

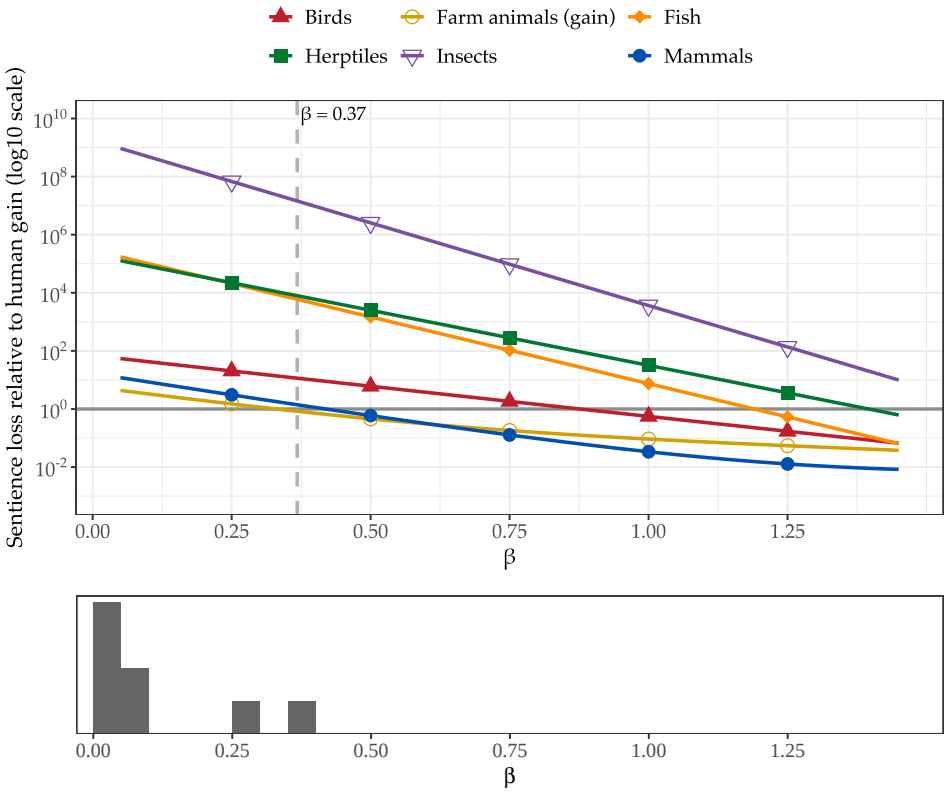

Notes: The y-axis is the sentience loss or gain for each taxon expressed relative to human sentience gains (as defined in (10)). The x-axis is the sentience-to-neurons elasticity β. The "Farm animals" category includes all farmed animals, irrespective of taxon. We use population data at the most granular level of disaggregation available.

of life due to economic growth in Section 5.5.[19]

We report our results for various values of β, the elasticity of sentience to neuron counts. Varying β simultaneously spans alternative cardinalizations of cognition and different choices of cognitive tests. Our baseline is conservative: we use the largest value of β across the A-not-B, cylinder and numeration tests, which cover the most species. This choice yields $\beta = \hat{\beta}^{\text{A-B}} = 0.37$.

5.2 The Aggregate Sentience Cost of Biodiversity Loss

We start with aggregate sentience gains or losses by animal taxa (10), as a function of β. Figure 4 displays our results. Sentience losses for every taxon exceed human sentience

[19] Alternatively, we can express sentience changes relative to the sentience of a single human:

$$\frac{-\Delta \mathcal{W}_{st}}{\overline{u}_0} = \frac{\overline{u}_s}{\overline{u}_0}\left(N_{s,1970} - N_{s,t}\right) \geq -\omega_s \ \left(N_{s,1970} - N_{s,t}\right).$$

gains at our baseline estimate $\beta = \hat{\beta}^{\text{A-B}}$. For birds alone, the sentience losses are one order of magnitude larger than human sentience gains. For herptiles and fish, sentience losses are four orders of magnitude larger than human sentience gains. Insect sentience losses are seven orders of magnitude larger. At larger values of β, these gaps fall. Yet, even when sentience scales linearly with neuron counts ($\beta = 1$), sentience losses for fish and herptiles remain larger than human sentience gains, and losses for insects remain several orders of magnitude larger than human sentience gains.

To assess whether $\beta = \hat{\beta}^{\text{A-B}}$ or $\beta = 1$ are plausible values, we additionally plot the histogram of estimates of β across all the cognitive tests that cover multiple taxonomic groups. Our baseline value $\beta = \hat{\beta}^{\text{A-B}}$ is already conservative: no other test implies a stronger scaling of cognition with neuron counts than the A-not-B test. In particular, choosing $\beta = 1$ would be much more conservative than what the data suggest.

Figure 4 also depicts sentience gains from rising farmed animal populations. This group represents a marginal adjustment to positive contributions to total sentience. While farmed animals represent an important proportion of earthborne biomass, their sentience remains below that of humans. Therefore, they represent at most a doubling of positive contributions to sentience. By contrast, wild animals represent a smaller proportion of biomass but substantial aggregate sentience because small animals can be strongly sentient. Because the sentience costs of biodiversity loss are orders of magnitude larger than human sentience gains, farmed animals make little difference to this comparison. If we further adjusted their contribution to sentience due to their arguably lower quality of life, their positive contribution to aggregate sentience would be even smaller.

Next, we aggregate sentience across animal taxa following equation (1). Figure 5 shows total animal sentience losses relative to human sentience gains. We consider four cases. First, we include all wild animal welfare losses. Second, we consider all wild animal welfare losses of vertebrates, but exclude insects whose sentience is less clearly established. Third, we include the combined welfare change of wild and farmed animals. Fourth, we include the combined welfare change of wild animals and farmed animals, excluding insects.

We find that animal sentience losses exceed human sentience gains by multiple orders of magnitude for all cases and across all plausible values for β. At our baseline estimate of $\beta = 0.37$, the sentience costs of biodiversity loss are more than 10,000 times larger than human sentience gains excluding invertebrates. When we include

FIGURE 5: Total Animal Sentience Loss Relative to Human Gain

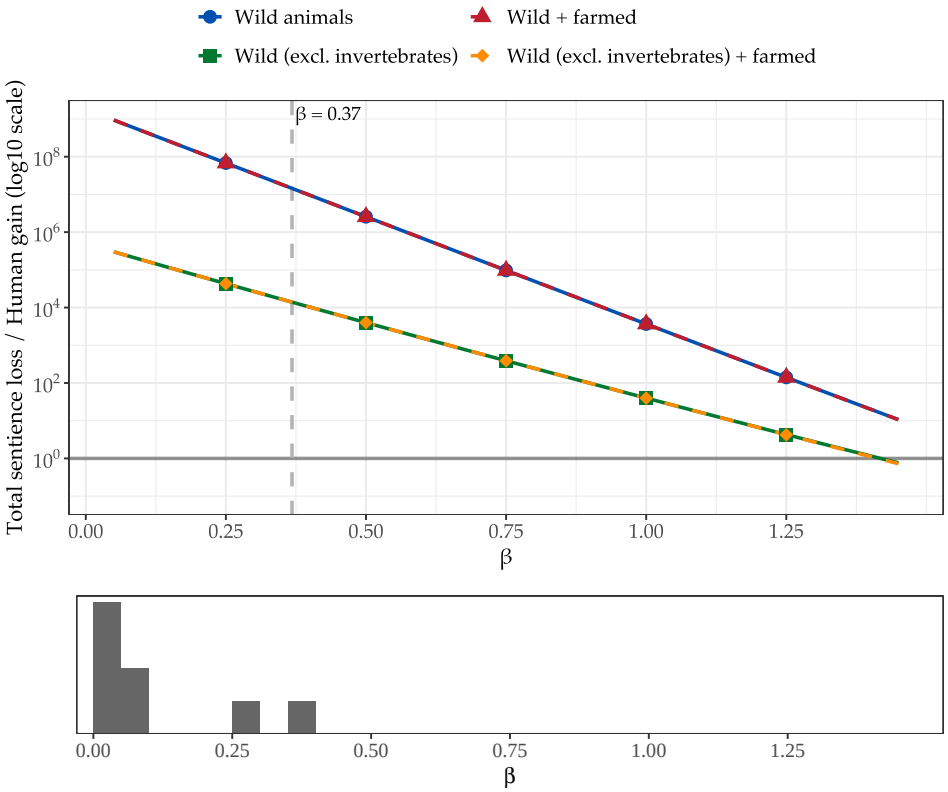

Notes: The y-axis is the sentience loss or gain for all taxa expressed relative to human sentience gain (as defined in (10)). The x-axis is the sentience-to-neurons elasticity β. The "Farm animals" category includes all farmed animals, irrespective of taxon. We use population data prioritizing the most granular level of disaggregation.

invertebrates, the ratio rises to more than ten million. At the conservative value of $\beta = 1$ (with less empirical support), the sentience cost of biodiversity loss remains more than 50 times larger than human sentience gains after excluding insects. Including insects makes this ratio exceed 1,000.

In the Appendix, we assess how our results change as we vary assumptions. First, we use more aggregated population data for mammals and birds, defining neuron counts for each taxon as either the mean or the minimum of observed neuron counts within that taxon (Figure E.1). Second, we use relative welfare capacities estimated from pallial neuron counts, as opposed to total neuron counts (Figure E.2). Third, we apply the relative fish population decline from the LPI to all fish as opposed to excluding mesopelagic fish (Figure E.3). The sentience costs of biodiversity loss exceed human welfare gains by multiple orders of magnitude across all these specifications.

5.3 Top Sentience Convexity

Our results rely on the estimated mapping between neuron counts and sentience. Figures 4, 5 and 7 already show how our results depend on the convexity β of this mapping. However, one might reasonably suspect that this mapping is not exactly log-linear. In particular, the relationship between neuron counts and cognitive abilities may steepen as we go from the smartest animals to humans.

Excess convexity at the top might be difficult to detect empirically: most cognitive tests are designed so that they might be completed by animals and humans alike. For instance, our cognitive tests do not include inverting a 3×3 matrix by hand, at which humans might succeed but most animals would likely fail (although humans usually require extensive prior training to even take this sort of test). These considerations suggest that our estimate of β might be biased downward at the top of the neuron count distribution.

To gauge whether this concern is likely to affect our conclusions, we leverage the Primate Cognition Test Battery. These tests are more demanding than the A-not-B and cylinder tests, and are primarily run on primates. Figure D.4 shows our estimates of the scaling between neuron counts and performance on the PCTB. Across the 16 tests, most slopes are below our baseline $\hat{\beta}^{\text{A-B}} = 0.37$, but two tests (social learning and tool use) deliver substantially higher slopes. We take the largest estimated slope $\beta_2 = \hat{\beta}^{\text{PCTB}} = 2.1$ among the primate cognition tests as our most conservative value, and use it to parametrize a piecewise relationship between cognition and neuron counts. We assume that it is governed by β_2 between humans and the top primate, and by β_1 between the top primate and other animals. This choice additionally downweights primates, and hence all animals, relative to humans using β_2. Namely, we define a new sentience value ω_s^{conv} by

$$\omega_s^{\text{conv}} = \begin{cases} \left(\frac{z_s}{z_{\text{top primate}}}\right)^{\beta_1} \left(\frac{z_{\text{top primate}}}{z_{\text{human}}}\right)^{\beta_2}, & \text{if } z_s \leq z_{\text{top primate}} \\ \left(\frac{z_s}{z_{\text{human}}}\right)^{\beta_2}, & \text{if } z_{\text{top primate}} < z_s \leq z_{\text{human}} \end{cases}$$

Figure 6 displays how our results jointly depend on β_1, β_2. Panel 6a shows that incorporating higher sentience convexity at the top (at $\beta_2 = \hat{\beta}^{\text{PCTB}} = 2.1$) reduces sentience losses for all groups by less than one order of magnitude. Mammal sentience losses are closest to human sentience gains before the adjustment, and thus fall below

FIGURE 6: Animal Sentience Loss Relative to Human Gain under ω^{conv}

(A) By Group, $\beta_1 = \hat{\beta}^{\text{A-B}}$ and by β_2

(B) Aggregate, $\beta_1 = \hat{\beta}^{\text{A-B}}$ and by β_2

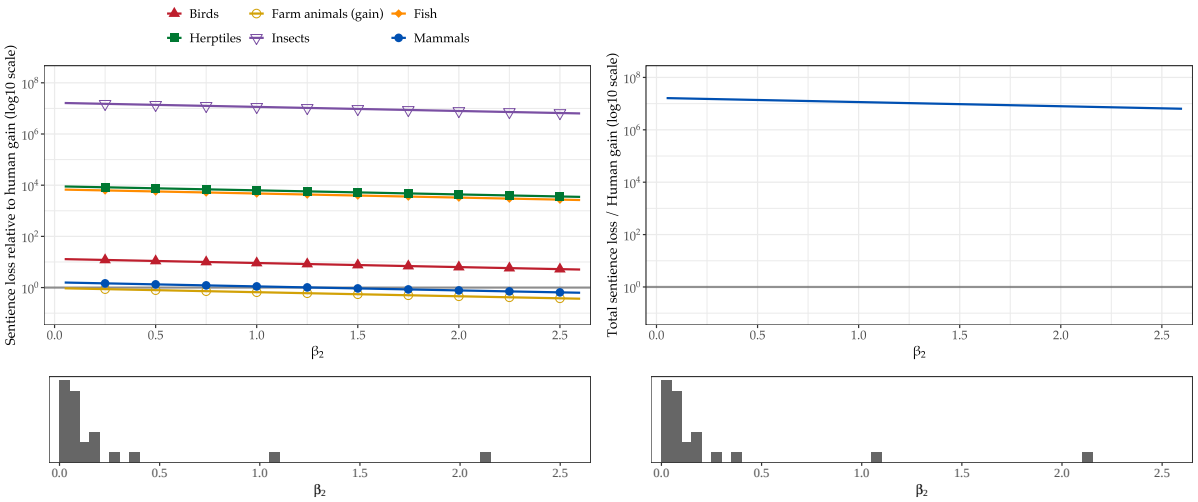

(C) By Group, $\beta_2 = \hat{\beta}^{\text{PCTB}}$ and by β_1

(D) Aggregate, $\beta_2 = \hat{\beta}^{\text{PCTB}}$ and by β_1

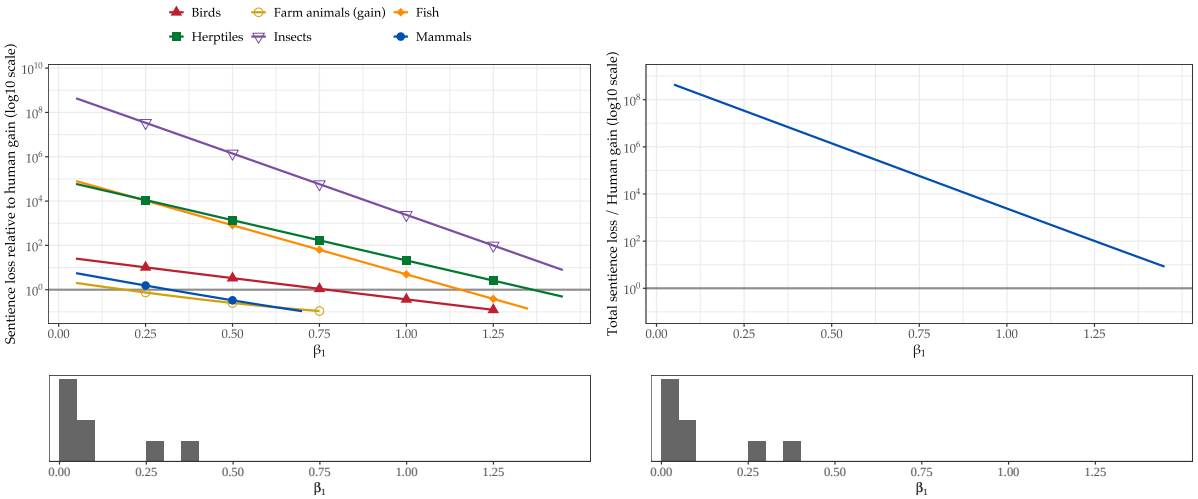

Notes: The y-axis is the sentience loss or gain for all taxa expressed relative to human sentience gain (as defined in (10)). The welfare loss is calculated using ω_s^{conv}. The x-axis is the sentience-to-neurons elasticity β. The "Farm animals" category includes all farmed animals, irrespective of taxon. Panel (a): animal sentience loss by group under $\beta_1 = \hat{\beta}^{\text{A-B}}$ as a function of β_2. Panel (b): aggregate animal sentience loss under $\beta_1 = \hat{\beta}^{\text{A-B}}$ as a function of β_2. Panel (c): animal sentience loss by group under $\beta_2 = \hat{\beta}^{\text{PCTB}}$ as a function of β_1. Panel (d): aggregate animal sentience loss under $\beta_2 = \hat{\beta}^{\text{PCTB}}$ as a function of β_1.

human sentience gains. Other taxonomic groups continue to exhibit sentience losses that are one or several orders of magnitude larger than human sentience gains. As a result, panel 6b shows that aggregate animal sentience losses remain multiple orders of magnitude above human sentience gains. Below the graph, we display the histogram of estimates of β across primate cognition tests, which highlights that our baseline value of $\hat{\beta}^{\text{PCTB}} = 2.1$ is strongly conservative. In addition, it is worth noting that the lines in panels 6a and 6b are relatively flat.

These results indicate that our conclusions are not driven by humans being relatively close to animals with the highest cognition. Rather, they reflect the fact that cognitive, and hence sentience, differences are relatively compressed across animal species.

Panel 6c displays animal sentience losses by group as a function of β_1, holding $\beta_2 = \hat{\beta}^{\text{PCTB}}$. After incorporating higher convexity at the top of the sentience distribution, and assuming a large value of β_1, the sentience loss of several animal groups becomes smaller than human sentience gains. Importantly, values of β_1 at which mammal and bird sentience losses become smaller than human gains are much higher than what the data suggest. Even then, panel 6d indicates that aggregate animal sentience losses are driven by more numerous groups such as fish and herptiles and continue to be multiple orders of magnitude above human sentience gains.

5.4 Animal Welfare Weight

We now convert the aggregate sentience costs from biodiversity loss into welfare. For this exercise, we consider animal welfare $\mathcal{W}_{At}(\boldsymbol{\lambda})$ as a function of the vector of animal welfare weights $\boldsymbol{\lambda}$. For simplicity, we use a common welfare weight $\lambda_s = \lambda \leq 1$ for all species other than humans, so $s > 0$. $\lambda < 1$ corresponds to a social welfare function that is more human-centric than what would be consistent with a strict interpretation of utilitarianism.

To ease the interpretation of our results, we convert our sentience metrics into a willingness to pay to preserve animal life. We consider three species groups or species for illustration: great apes, crows and salmon. For this calculation, we need two objects: a Value of Statistical Life (VSL) for humans, and life expectancy for each species.[20]

We start from a standard human VSL of USD 10 million. We scale this human value of a statistical life by two factors: expected remaining life-years for each species relative to 40 years for humans, and relative sentience. Great apes are assigned 20 remaining years from adulthood, crows are assigned 6 remaining years, and salmon 3 remaining years.

With a common welfare weight $\lambda = 1$, a great ape's life is valued at USD 3.2 million, a crow's life at USD 180,000 and a salmon's life at USD 2,100. These values exceed stated preferences for biodiversity protection by several orders of magnitude (Loomis and White 1996). Next, we consider a welfare weight $\lambda = 10^{-4}$ that approximately equalizes

[20] This calculation requires us to adjust for life expectancy because we consider the value of a single life rather than a change in steady-state population.

animal welfare losses and human welfare gains under conservative assumptions (when excluding insects). In that case, a great ape's life is valued at 320 dollars, a crow's life at 18 dollars and a salmon's life at 21 cents.

Two conclusions emerge from these calculations. First, the willingness to pay for any given animal life depends crucially on the welfare weight of animals. While settling the exact value is beyond the scope of this paper, the magnitude of the potential welfare costs of biodiversity losses highlights its importance. Second, in the aggregate, even at small values of animal welfare weights, the welfare costs of biodiversity loss are at least as large as historical welfare gains to humans. Even at a smaller willingness to pay for any individual animal life, the scope of population declines is such that overall welfare costs remain large.

Importantly, even breaking even with human gains does not preclude large aggregate dollar-equivalent losses. We define human gains to include the contribution of population gains, which thus exceed annual Gross Domestic Product (GDP) at conventional VSLs. At an animal welfare weight of 10^{-4} that approximately equalizes animal losses excluding insects and human gains, the welfare cost from biodiversity loss represents 112% of global GDP in 2025.

5.5 Quality of Life and Economic Growth

Quality of life improvements from growth. So far, we have assumed that all species' quality of life is fixed at their state of nature. In practice, human quality of life has improved thanks to economic growth, so that $\tilde{u}_{0t} > \tilde{u}_{00}$ for humans. We now incorporate the benefits of economic growth for humans while assuming that animals remained in the state of nature. In this exercise, we first consider an animal welfare weight of $\lambda = 1$ for comparability, so animal welfare losses coincide with animal sentience losses.

To parametrize the size of the welfare gain inclusive of growth, we consider the following utility function for humans, consistent with the definition of welfare ranges:

$$u_0(C) = \begin{cases} \overline{u}_0 + \log \frac{C}{\overline{C}} & \text{if } C \leq \overline{C} \\ \overline{u}_0 & \text{otherwise,} \end{cases} \tag{11}$$

where C is consumption per capita, and $\overline{C}$ denotes the value of consumption beyond which there are no utility gains anymore. The parameter $\overline{C}$ is necessary to remain consistent with the assumption that utility is bounded. In practice, we take $\overline{C}$ to be

large enough so that this bound is never binding up to 2026. The functional form (11) is similar to Hall and Jones (2007) and Jones and Klenow (2016).

The utility function (11) requires us to estimate $\overline{u}_0$: it controls the marginal rate of substitution between consumption and life. We leverage that $\overline{u}_0$ is related to the value of statistical life (VSL). In particular:[21]

$$\overline{u}_0 + \log \frac{C_t}{\overline{C}} = \frac{\text{VSL}_t}{C_t}. \tag{12}$$

We may thus estimate a combination of $\overline{u}_0$ and $\overline{C}$ given a VSL and consumption.

The VSL in the United States is USD 10 million in 2025. As our model is static, we convert the VSL to a flow measure by dividing it by the average number of years left to live for the average adult, which we take to be 40. Consumption per capita is USD 62,387 in the United States in 2025. We further adjust for the gap between consumption in the United States and the world average. We obtain:

$$\overline{u}_0 + \log \frac{C^{\text{World}}_{2025}}{\overline{C}} = \overline{u}_0 + \log \frac{C^{\text{US}}_{2025}}{\overline{C}} + \log \frac{C^{\text{World}}_{2025}}{C^{\text{US}}_{2025}} = \frac{\text{VSL}^{\text{US}}_{2025}}{C^{\text{US}}_{2025}} + \log \frac{C^{\text{World}}_{2025}}{C^{\text{US}}_{2025}} = 2.17.$$

With this estimate in hand, we are ready to construct an upper bound on the welfare change for humans inclusive of economic growth that does not depend on $\overline{C}$.

Proposition 2. *(Upper bound on welfare changes with growth) Suppose that $v_t = \overline{u}_0 + \log \frac{C_t}{\overline{C}}$ is given, and that human population and consumption increase weakly between date τ and date t: $C_\tau \leq C_t$ and $N_{0,\tau} \leq N_{0,t}$. The human welfare change between date τ and date t satisfies:*

$$\frac{\Delta \mathcal{W}_{0t}}{\overline{u}_0} \leq \Delta N_{0,t} + \frac{N_{0,\tau}}{v_t} \log \frac{C_t}{C_\tau}.$$

Proof. Start from: $\Delta \mathcal{W}_{0t} = N_{0,t} \left(\overline{u}_0 + \log \frac{C_t}{\overline{C}} \right) - N_{0,\tau} \left(\overline{u}_0 + \log \frac{C_\tau}{\overline{C}} \right)$. Rearrange to obtain: $\frac{\Delta \mathcal{W}_{0t}}{\overline{u}_0} = \frac{v_t}{\overline{u}_0} \Delta N_{0,t} + \frac{N_{0,\tau}}{\overline{u}_0} \log \frac{C_t}{C_\tau}$. To conclude the proof, observe that $\frac{v_t}{\overline{u}_0} = \frac{\overline{u}_0 + \log \frac{C_t}{\overline{C}}}{\overline{u}_0} \leq 1$ since $C_t \leq \overline{C}$, and that $\frac{1}{\overline{u}_0} \leq \frac{1}{v_t}$. □

The first component of the expression in Proposition 2 is identical to the expression without growth. This first component is an upper bound because consumption might be below the bliss point. The second component encodes the welfare gain from growth. It scales initial population only because consumption gains by population in the final

[21] Formally, we extend $u_0(C, p) = p(\overline{u}_0 + \log C/\overline{C})$, where p denotes the probability of surviving. Then: $\text{VSL}_t = \frac{\partial u(C_t)}{\partial p}\Big|_{p=1} \Big/ \frac{\partial u(C_t)}{\partial C}$. Then: $\text{VSL}_t = \left(\overline{u}_0 + \log \frac{C_t}{\overline{C}} \right) C_t$. Rearranging delivers (12).

FIGURE 7: Total Animal Welfare Loss Relative to Human Welfare Gain With Growth

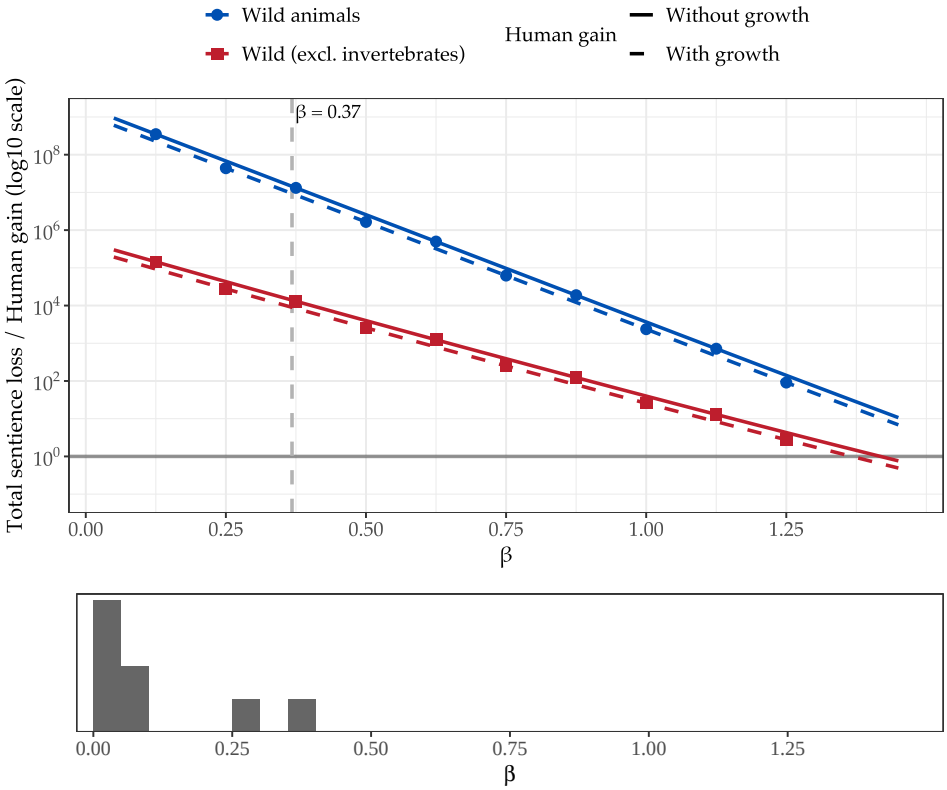

Notes: The y-axis is the welfare loss or gain for all taxa expressed relative to human gain (as defined in (10)), where the human welfare gain either excludes growth (solid lines) or includes growth (dashed lines). The x-axis is the sentience-to-neurons elasticity β. The "Farm animals" category includes all farmed animals, irrespective of taxon. We use population data prioritizing the most granular level of disaggregation.

period are already bounded by the first component. The second component is bounded above by the empirical estimate $1/v_t$ of the inverse of the welfare range. Proposition 2 resembles the calculations in Adhami et al. (2023) who use a total utilitarian welfare criterion to incorporate population growth into growth accounting.

We use Proposition 2 to bound the welfare gains from economic growth. Figure 7 depicts aggregate animal sentience losses relative to human welfare gains.

We find that adjusting for economic growth does not affect our conclusions meaningfully. To understand why, observe that human welfare gains $N_{0,2026} - N_{0,1970}$ relative to human sentience between 1970 and 2026 absent economic growth represent the population increase of 4.5 billion human lives. The welfare gains from economic growth add 2.5 billion human-life equivalents. This adjustment is of the same order of magnitude as human welfare gains from expanding population, consistent with Jones and Klenow (2016). Since the aggregate sentience costs from biodiversity loss are at least 10,000 times larger than human sentience gains, quality of life improvements due to economic

growth leave our conclusions largely unchanged.

Quality of life in the state of nature. So far, we have assumed that humans and animals have equal quality of life in a state of nature, and that the only reason for divergence in quality of life is economic growth. There are two potential caveats to this approach.

First, humans and other animals may have different qualities of life in the state of nature. Which direction this might go is not obvious. Human cognition and reflection could make the state of nature more or less enjoyable than for an animal. To be conservative, we assume that animals have a lower quality of life in the state of nature. This assumption is isomorphic to assigning a smaller welfare weight $\lambda_s < 1$ to animals. Given that animal sentience losses are 10,000 to 10^7 times larger than human gains, one would need to assume that animals achieve only 10^{-7} to 10^{-4} of their welfare range relative to humans to break even.

We can also relax the normalization that utility is weakly positive introduced in Section 2.3 by choosing $\lambda_s < 0$. Doing so introduces the idea that animals would prefer not to live in the state of nature and assigns negative utility to them. If a sizable share of wild animals have negative utility in the state of nature, our conclusions would change meaningfully: indeed, removing net negative lives is a net positive.

Second, there may have been diverging trends in human and animal quality of life since the state of nature that go beyond the adjustment for economic growth. One candidate is declining wild animal quality of life due to habitat loss and other environmental pressures. A second candidate is farm animals' quality of life being substantially below that experienced in the state of nature. Accounting for these would only increase our estimated welfare losses, so that omitting these forces is conservative.[22]

6. Conclusion

This paper explores the welfare costs of the loss of animal life in a consequentialist total utilitarian framework that incorporates the intrinsic value of living beings through their sentience. Using cognitive tests to construct a conservative lower bound on sentience under plausible neuroscientific assumptions, we estimate the relationship between cognition and neuron counts to cardinalize sentience. Combined with population

[22] Even negative utility for farmed animals would not affect our results much because their aggregate sentience is several orders of magnitude below that of wild animals.

and neuron data for a wide range of animal species, we find that aggregate sentience losses of animals in the past half century outstrip the sentience gains to humans from population growth by at least 10,000-fold. Breaking even in welfare terms requires assigning a correspondingly small welfare weight to animals.

These results reveal a large discrepancy between a sentience-based valuation of animal welfare and humans' valuation of animal welfare as revealed by their behavior and willingness to pay. Through the lens of our framework, this suggests that humans assign welfare weights to other animals that are substantially below one. We do not argue for a particular value of the animal welfare weight. Rather, our objective is to highlight the importance of this question given the substantial losses in animal sentience over the past half century. In particular, seemingly small departures from a zero animal welfare weight can lead to large aggregate welfare costs from biodiversity loss.

REFERENCES

Adhami, Mohamad, Mark Bils, Charles I Jones, and Peter J Klenow, "Population and Welfare: Measuring Growth when Life is Worth Living," Working Paper 31999, National Bureau of Economic Research December 2023.

Ahrens, Misha B., Michael B. Orger, Drew N. Robson, Jennifer M. Li, and Philipp J. Keller, "Whole-brain functional imaging at cellular resolution using light-sheet microscopy," *Nature Methods*, May 2013, *10* (5), 413–420.

Bar-On, Yinon M., Rob Phillips, and Ron Milo, "The biomass distribution on Earth," *Proceedings of the National Academy of Sciences*, 2018, *115* (25), 6506–6511.

Barbier, Edward B, "Valuing ecosystem services as productive inputs," *Economic policy*, 2007, *22* (49), 178–229.

Barron, Andrew B. and Faelan Mourmourakis, "The Relationship between Cognition and Brain Size or Neuron Number," *Brain Behavior and Evolution*, July 2023, *99* (2), 109–122.

Bateman, Ian J, Amii R Harwood, Georgina M Mace, Robert T Watson, David J Abson, Barnaby Andrews, Amy Binner, Andrew Crowe, Brett H Day, Steve Dugdale et al., "Bringing ecosystem services into economic decision-making: land use in the United Kingdom," *science*, 2013, *341* (6141), 45–50.

Birch, Jonathan, *The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI*, Oxford: Oxford University Press, 2024.

Blackorby, Charles and David Donaldson, "Pigs and Guinea Pigs: A Note on the Ethics of

Animal Exploitation," *The Economic Journal*, 1992, *102* (415), 1345–1369.

Browning, Heather, "If I Could Talk to the Animals: Measuring Subjective Animal Welfare." PhD dissertation, Australian National University, Canberra 2019. Page reference: p. 153.

___ **and Walter Veit**, "Positive Wild Animal Welfare," *Biology & Philosophy*, 2023, *38* (6), 48.

Budolfson, Mark, Romain Espinosa, Bob Fischer, and Nicolas Treich, "Monetizing Animal Welfare Impacts for Benefit–Cost Analysis," *Journal of Benefit-Cost Analysis*, 2024, *15* (S1), 206–223.

Carlier, Alexis and Nicolas Treich, "Directly valuing animal welfare in (environmental) economics," *International Review of Environmental and Resource Economics*, 2020, *14* (1), 113–152.

Chalmers, David J., "Facing Up to the Problem of Consciousness," *Journal of Consciousness Studies*, 1995, *2* (3), 200–219.

Costanza, Robert, Ralph d'Arge, Rudolf De Groot, Stephen Farber, Monica Grasso, Bruce Hannon, Karin Limburg, Shahid Naeem, Robert V O'neill, Jose Paruelo et al., "The value of the world's ecosystem services and natural capital," *Nature*, 1997, *387* (6630), 253–260.

Crisp, Roger, *Reasons and the Good*, Oxford: Clarendon Press, 2006.

Daily, Gretchen C., Tore Söderqvist, Sara Aniyar, Kenneth Arrow, Partha Dasgupta, Paul R. Ehrlich, Carl Folke, AnnMari Jansson, Bengt-Owe Jansson, Nils Kautsky, Simon Levin, Jane Lubchenco, Karl-Göran Mäler, David Simpson, David Starrett, David Tilman, and Brian Walker, "The Value of Nature and the Nature of Value," *Science*, 2000, *289* (5478), 395–396.

Damasio, Antonio, *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*, New York: Harcourt, 1999.

___ **and Gil B. Carvalho**, "The Nature of Feelings: Evolutionary and Neurobiological Origins," *Nature Reviews Neuroscience*, 2013, *14* (2), 143–152.

DeGrazia, David, *Taking Animals Seriously: Mental Life and Moral Status*, Cambridge: Cambridge University Press, 1996.

Dehaene, Stanislas, Hakwan Lau, and Sid Kouider, "What Is Consciousness, and Could Machines Have It?," *Science*, 2017, *358* (6362), 486–492.

Dicke, Ursula and Gerhard Roth, "Neuronal Factors Determining High Intelligence," *Philosophical Transactions of the Royal Society B*, 2016, *371* (1685), 20150180.

Dirzo, Rodolfo, Hillary S. Young, Mauro Galetti, Gerardo Ceballos, Nick J. B. Isaac, and Ben Collen, "Defaunation in the Anthropocene," *Science*, 2014, *345* (6195), 401–406.

Eden, Maya, Jojo Lee, and Nicolas Treich, "Global Welfare Trends," *Master Thesis*, 2026.

Espinosa, Romain and Nicolas Treich, "Animal welfare: antispeciesism, veganism and a" life worth living"," *Social Choice and Welfare*, 2021, pp. 531–548.

___ **and** ___ , "The animal-welfare levy," *Journal of the European Economic Association*, 2025, p. jvaf063.

Feldman, Fred, "The Good Life: A Defense of Attitudinal Hedonism," *Philosophy and Phenomenological Research*, 2002, *65* (3), 604–628.

___ , *Pleasure and the Good Life*, Oxford: Clarendon Press, 2004.

Fleurbaey, Marc and Martin Van der Linden, "Fair Social Ordering, Egalitarianism, and Animal Welfare," *American Economic Journal: Microeconomics*, November 2021, *13* (4), 466–91.

Frank, Eyal and Anant Sudarshan, "The social costs of keystone species collapse: evidence from the decline of vultures in India," *American Economic Review*, 2024, *114* (10), 3007–3040.

Frank, Eyal G, "The economic impacts of ecosystem disruptions: Costs from substituting biological pest control," *Science*, 2024, *385* (6713), eadg0344.

Froese, R. and D. Pauly, eds, "FishBase," www.fishbase.org 2026. Version (02/2026).

Giglio, Stefano, Theresa Kuchler, Johannes Stroebel, and Olivier Wang, "Nature and Biodiversity Loss: A Research Agenda for Financial Economics," Working Paper 34286, National Bureau of Economic Research September 2025.

Ginsburg, Simona and Eva Jablonka, *The Evolution of the Sensitive Soul: Learning and the Origins of Consciousness*, Cambridge, MA: MIT Press, 2019.

Godfrey-Smith, Peter, "Animal Evolution and the Origins of Experience," in David Livingstone Smith, ed., *How Biology Shapes Philosophy: New Foundations for Naturalism*, Cambridge: Cambridge University Press, 2017, pp. 51–71.

Greenspoon, Lior, Noam Ramot, Uri Moran, Uri Roll, Rob Phillips, Elad Noor, and Ron Milo, "The global biomass of mammals since 1850," *Nature Communications*, 2025, *16*, 8338.

Groff, Zach and Yew-Kwang Ng, "Does Suffering Dominate Enjoyment in the Animal Kingdom? An Update to Welfare Biology," *Biology & Philosophy*, 2019, *34* (2), 16.

Hall, Robert E. and Charles I. Jones, "The Value of Life and the Rise in Health Spending," *The Quarterly Journal of Economics*, 02 2007, *122* (1), 39–72.

Hare, R. M., *Moral Thinking: Its Levels, Method, and Point*, Oxford: Clarendon Press, 1981.

Herculano-Houzel, Suzana, "Not All Brains Are Made the Same: New Views on Brain Scaling in Evolution," *Brain Behavior and Evolution*, 2011, *78* (1), 22–36.

___ , "Numbers of neurons as biological correlates of cognitive capability," *Current Opinion in Behavioral Sciences*, August 2017, *16*, 1–7.

___ **, Bruno Mota, and Roberto Lent**, "Cellular scaling rules for rodent brains," *Proceedings of the National Academy of Sciences*, 2006, *103* (32), 12138–12143.

___ **, Christine E. Collins, Peiyan Wong, and Jon H. Kaas**, "Cellular scaling rules for primate brains," *Proceedings of the National Academy of Sciences*, 2007, *104* (9), 3562–3567.

___ **, Kenneth Catania, Paul R. Manger, and Jon H. Kaas**, "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*, September 2015, *86* (3-4),

145–163.

Johansson-Stenman, Olof, "Animal Welfare and Social Decisions: Is It Time to Take Bentham Seriously?," *Ecological Economics*, 2018, *145*, 90–103.

Jones, Charles I. and Peter J. Klenow, "Beyond GDP? Welfare across Countries and Time," *American Economic Review*, September 2016, *106* (9), 2426–57.

Jones, Kate E., Jon Bielby, Marcel Cardillo, Susanne A. Fritz, Justin O'Dell, C. David L. Orme, Kamran Safi, Wes Sechrest, Elizabeth H. Boakes, Chris Carbone, Christina Connolly, Michael J. Cutts, Janine K. Foster, Richard Grenyer, Michael Habib, Christopher A. Plaster, Samantha A. Price, Elizabeth A. Rigby, Janna Rist, Amber Teacher, Olaf R. P. Bininda-Emonds, John L. Gittleman, Georgina M. Mace, and Andy Purvis, "PanTHERIA: a species-level database of life history, ecology, and geography of extant and recently extinct mammals," *Ecology*, 2009, *90* (9), 2648–2648. _eprint: https://esajournals.onlinelibrary.wiley.com/doi/pdf/10.1890/08-1494.1.

Kim, Jaegwon, *Physicalism, or Something Near Enough*, Princeton: Princeton University Press, 2005.

Koob, George F. and Nora D. Volkow, "Neurocircuitry of Addiction," *Neuropsychopharmacology*, 2010, *35* (1), 217–238.

Krasheninnikova, Anastasia, Roberta Berardi, Mari-Ann Lind, Laurie O'Neill, and Auguste M.P. von Bayern, "Primate cognition test battery in parrots," *Behaviour*, 2019, *156* (5/8), 721–761.

Krutilla, John V., "Conservation Reconsidered," *American Economic Review*, 1967, *57* (4), 777–786.

Kuruc, Kevin, Mark Budolfson, and Dean Spears, "Population Axioms and Population Ethics," 2022. Survey paper.

Kverková, Kristina, Lucie Marhounová, Alexandra Polonyiová, Martin Kocourek, Yicheng Zhang, Seweryn Olkowicz, Barbora Straková, Zuzana Pavelková, Roman Vodička, Daniel Frynta, and Pavel Němec, "The evolution of brain neuron numbers in amniotes," *Proceedings of the National Academy of Sciences*, March 2022, *119* (11), e2121624119.

Lamme, Victor A. F., "Towards a True Neural Stance on Consciousness," *Trends in Cognitive Sciences*, 2006, *10* (11), 494–501.

LeDoux, Joseph E. and Richard Brown, "A Higher-Order Theory of Emotional Consciousness," *Proceedings of the National Academy of Sciences*, 2017, *114* (10), E2016–E2025.

Levy, Dino J. and Paul W. Glimcher, "The Root of All Value: A Neural Common Currency for Choice," *Current Opinion in Neurobiology*, 2012, *22* (6), 1027–1038.

Loomis, John B. and Douglas S. White, "Economic benefits of rare and endangered species: summary and meta-analysis," *Ecological Economics*, 1996, *18* (3), 197–206.

MacLean, Evan L., Brian Hare, Charles L. Nunn, Elsa Addessi, Federica Amici, Rindy C. Anderson, Filippo Aureli, Joseph M. Baker, Amanda E. Bania, Allison M. Barnard, Neeltje J.

Boogert, Elizabeth M. Brannon, Emily E. Bray, Joel Bray, Lauren J. N. Brent, Judith M. Burkart, Josep Call, Jessica F. Cantlon, Lucy G. Cheke, Nicola S. Clayton, Mikel M. Delgado, Louis J. DiVincenti, Kazuo Fujita, Esther Herrmann, Chihiro Hiramatsu, Lucia F. Jacobs, Kerry E. Jordan, Jennifer R. Laude, Kristin L. Leimgruber, Emily J. E. Messer, Antonio C. de A. Moura, Ljerka Ostojić, Alejandra Picard, Michael L. Platt, Joshua M. Plotnik, Friederike Range, Simon M. Reader, Rachna B. Reddy, Aaron A. Sandel, Laurie R. Santos, Katrin Schumann, Amanda M. Seed, Kendra B. Sewall, Rachael C. Shaw, Katie E. Slocombe, Yanjie Su, Ayaka Takimoto, Jingzhi Tan, Ruoting Tao, Carel P. van Schaik, Zsófia Virányi, Elisabetta Visalberghi, Jordan C. Wade, Arii Watanabe, Jane Widness, Julie K. Young, Thomas R. Zentall, and Yini Zhao, "The evolution of self-control," *Proceedings of the National Academy of Sciences*, May 2014, *111* (20), E2140–E2148.

Marhounová, Lucie, Alexander Kotrschal, Kristina Kverková, Niclas Kolm, and Pavel Němec, "Artificial selection on brain size leads to matching changes in overall number of neurons," *Evolution*, September 2019, *73* (9), 2003–2012.

Merker, Bjorn, "Consciousness without a Cerebral Cortex: A Challenge for Neuroscience and Medicine," *Behavioral and Brain Sciences*, 2007, *30* (1), 63–81.

Metrick, Andrew and Martin L Weitzman, "Conflicts and choices in biodiversity preservation," *Journal of Economic Perspectives*, 1998, *12* (3), 21–34.

Olkowicz, Seweryn, Martin Kocourek, Radek K. Lučan, Michal Porteš, W. Tecumseh Fitch, Suzana Herculano-Houzel, and Pavel Němec, "Birds have primate-like numbers of neurons in the forebrain," *Proceedings of the National Academy of Sciences*, June 2016, *113* (26), 7255–7260.

Panksepp, Jaak, *Affective Neuroscience: The Foundations of Human and Animal Emotions*, New York: Oxford University Press, 1998.

Papineau, David, *Thinking About Consciousness*, Oxford: Oxford University Press, 2002.

Rausser, Gordon C and Arthur A Small, "Valuing research leads: bioprospecting and the conservation of genetic resources," *Journal of Political Economy*, 2000, *108* (1), 173–206.

Rigotti, Mattia, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi, "The Importance of Mixed Selectivity in Complex Cognitive Tasks," *Nature*, 2013, *497* (7451), 585–590.

Rosenthal, David M., *Consciousness and Mind*, Oxford: Oxford University Press, 2005.

Sarko, Diana K., Kenneth C. Catania, Duncan B. Leitch, Jon H. Kaas, and Suzana Herculano-Houzel, "Cellular scaling rules of insectivore brains," *Frontiers in Neuroanatomy*, 2009, *3*.

Simpson, R David, Roger A Sedjo, and John W Reid, "Valuing biodiversity for use in pharmaceutical research," *Journal of Political Economy*, 1996, *104* (1), 163–185.

Singer, Peter, *Animal Liberation*, New York: HarperCollins, 1975.

___ , *Practical Ethics*, Cambridge: Cambridge University Press, 1979.

Smith, Jane A and Kenneth M Boyd, "Lives in the balance: The ethics of using animals in biomedical research: The report of a working party of the institute of medical ethics," 1991.

Smithsonian Institution, "Numbers of Insects (Species and Individuals)," 1996.

Sol, Daniel, Seweryn Olkowicz, Ferran Sayol, Martin Kocourek, Yicheng Zhang, Lucie Marhounová, Christin Osadnik, Eva Corssmit, Joan Garcia-Porta, Thomas E. Martin, Louis Lefebvre, and Pavel Němec, "Neuron numbers link innovativeness with both absolute and relative brain size in birds," *Nature Ecology & Evolution*, September 2022, *6* (9), 1381–1389.

Stringer, Carsen, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D. Harris, "High-Dimensional Geometry of Population Responses in Visual Cortex," *Nature*, 2019, *571* (7765), 361–365.

Taylor, Charles A and Hannah Druckenmiller, "Wetlands, flooding, and the clean water act," *American Economic Review*, 2022, *112* (4), 1334–1363.

Teles, Magda C., Gonçalo M. Melo, Suzana Herculano-Houzel, and Rui F. Oliveira, "Larger Fish Have Larger Brains With More Neurons Across but Not Within Cohorts Raised in Different Growth Conditions," *Journal of Comparative Neurology*, 2025, *533* (9), e70090.

Tomasik, Brian, "How Many Wild Animals Are There?," Essays on Reducing Suffering 2018. [Online; accessed 2026-05-20].

Weitzman, Martin L, "The Noah's ark problem," *Econometrica*, 1998, pp. 1279–1298.

Wilman, Hamish, Jonathan Belmaker, Jennifer Simpson, Carolina de la Rosa, Marcelo M. Rivadeneira, and Walter Jetz, "EltonTraits 1.0: Species-level foraging attributes of the world's birds and mammals," *Ecology*, 2014, *95* (7), 2027–2027. _eprint: https://esajournals.onlinelibrary.wiley.com/doi/pdf/10.1890/13-1917.1.

Appendix

A. Definitions of Welfare

A.1 Utilitarianism: The Mental-State and Preference-Satisfaction Views

Utilitarianism is based on two main theories of welfare. The first theory is the mental-state theory (or more restrictively, the hedonic theory): welfare consists of positive mental states. The second theory is the preference-satisfaction theory: welfare consists of the satisfaction of preferences.[23] These theories coincide on the aspects that are relevant for our purposes in this paper. While it is beyond the scope of this paper to resolve the specific distinction between the mental-state and preference-satisfaction theories, we present the key arguments for and against the mental-state view, discuss the extent to which it is compatible with the preference-satisfaction theory, and explain why we adopt it for our project.

There are two key arguments in favor of the mental-state view. The first is subject-relativity. Welfare must be for the subject: it must matter to them. Only experiential states qualify without further argument. Crisp (2006) writes: "how could anything else benefit me except in so far as I enjoy it?" In their most general interpretation, preference-satisfaction theories fail this criterion: preference-satisfaction can imply positive welfare from the satisfaction of a preference that is entirely other-regarding and that the subject is unaware of. Of course, the standard counter-argument restricts the relevant preferences to self-regarding ones. With this restriction, however, defining "self-regarding" requires a prior account of what is good for the self, which is what a fundamental theory of welfare is supposed to provide. Under this counter-argument, the preference-satisfaction theory becomes very close to the mental-state theory rather than a distinct theory.

The second argument in favor of the mental-state theory is self-evidence. The intrinsic goodness of pleasant experience and the intrinsic badness of suffering are plausibly self-evident from the perspective of any individual. Sidgwick (1907) argued that one cannot intelligibly ask of pleasure "why is that good for me?", as the question

[23] The main alternatives outside utilitarianism are objective list theories and capability theories of welfare (Sen, Nussbaum, perfectionism), which treat goods such as autonomy, wisdom, and friendship as incommensurable and intrinsically valuable independently of their experiential impact.

is otherwise circular. The preference-satisfaction view, by contrast, does face the further question: "Why does the satisfaction of this or that preference make me better off?" Most, if not all, philosophers agree that some preferences should not count (uninformed, manipulated, antisocial), so the preference-satisfaction view must specify an additional restriction. The most common restriction is idealization: welfare is the satisfaction of those preferences the individual would have if they were well-informed and rational. But the criterion for "well-informed and rational" cannot itself be derived from preferences without circularity. The criterion has to be supplied by a substantive theory of welfare that specifies which preferences track welfare and which do not. The criterion thus relies on theory-laden judgments that preference satisfaction is meant to dispense with. As Rosati (1995) writes, the idealized agent whose preferences now define welfare may bear so little resemblance to the actual person that it is no longer clear whose welfare is being tracked.

Standard economics is sometimes read as endorsing a preference-satisfaction view of welfare. We emphasize that standard economics is entirely compatible with a mental-state view of welfare, because preference-based views can be read as evidential. The evidential reading states that preference-satisfaction is evidence of welfare, but welfare is something else. This reading is compatible with the mental-state view because welfare can be valenced experience, and preference-satisfaction reveals welfare.[24] We adopt the evidential reading, which also solves the issues with preference-satisfaction highlighted above.

While most work in economics does not explicitly distinguish the evidential reading from other interpretations, the evidential reading is consistent with economists' actual practice. For example, when economists use revealed preference to measure utility, they are not treating preferences as identical to welfare. Rather, they are treating them as indicators of welfare. This evidential reading is also in line with the behavioral economics literature, which distinguishes between *decision utility* (the weight of an outcome in choice) and *experienced utility* (the hedonic quality of an outcome). More recently, advances in behavioral economics and neuroeconomics have provided further support for the mental-state view by providing evidence on the neural basis of choice and utility (Caplin et al. 2010).

Specific interpretations of preference-satisfaction utilitarianism hold that genuine

[24]Another reading of preference-satisfaction views is the constitutive reading. In that interpretation, welfare is inherently preference-satisfaction. This reading is incompatible with the mental-state view, since preferences are not themselves valenced experiences.

preferences require language. Under this view, non-linguistic animals have no genuine interests and accordingly no moral status (Frey 1980). However, this criterion would also exclude human infants and severely cognitively impaired humans, which is at odds with current legal frameworks.

A.2 Alternatives to Utilitarianism

Theories consistent with sentientism. For Korsgaard's neo-Kantianism, the obligation is to treat all subjects as ends in themselves rather than as mere means. Respecting a being as an end in itself amounts to treating as valuable the ends that the being itself values. The sense in which an animal can be said to "value" an end is precisely that it has positively valenced experiences in pursuit of it.

For Nussbaum's capabilities approach, the imperative is to enable each being capable of "significant striving" to achieve a flourishing appropriate to its form of life. Nussbaum restricts the class of beings who can be wronged to those with "a felt orientation towards what is seen as good and a felt aversion to what is seen as bad" (Nussbaum 2023, p. 119), that is, those with valenced experience.

Theories inconsistent with sentientism. At one extreme, rationality-based moral status (orthodox Kantianism) contradicts sentientism, but it faces strong arguments against it. According to rationality-based moral status, moral status is tied to rational agency. Yet it cannot consistently include human infants and severely cognitively impaired humans without also including many sentient non-human animals.

At the other extreme, biocentrism or ecocentrism extends moral status to all living beings or to ecological wholes. However, it lacks a principled link between the entity (e.g., a plant or an ecosystem) and any subject for whom things can be non-instrumentally good or bad.

B. Theories of the Neural Basis of Valenced Experience

B.1 The Materialist View of Consciousness

Materialism and dualism. *Materialism* (or physicalism): conscious experience is fundamentally a physical phenomenon, ultimately realized in physical processes occurring

in the brain. To explain conscious experience, we do not need to posit any fundamentally non-physical properties, processes, forces, entities, substances, or laws.

Dualism: conscious experience, while dependent on brain function, is not itself a fundamentally physical phenomenon. To explain it, we do need to posit at least some fundamentally non-physical properties, processes, forces, entities, substances, or laws.

Materialism as the mainstream view. Materialism is the majority view among contemporary philosophers (Bourget and Chalmers 2014) and the near-universal view among neuroscientists (Koch et al. 2016).

The standard contemporary argument for materialism is the *causal-closure argument*, reconstructed most influentially by Papineau (2002) and Kim (2005). The argument has three premises. First, some conscious experiences have physical effects: experiencing pain causes one to withdraw, experiencing fear causes one to flee, and so on (mental causation). Second, all physical effects have fully physical causes, the principle of causal closure of the physical, supported by the cumulative success of physics in identifying complete physical-cause chains for physical effects. Third, the physical effects of conscious experiences are not, in general, overdetermined by distinct mental and physical causes. From these three premises it follows that at least some conscious experiences are themselves purely physical causes, and if that is so, materialism is true (at least for those conscious experiences that have physical effects).

The empirical balance has reinforced this argument. Over a century of experimental physiology and neuroscience has found zero evidence for a fundamentally non-physical force or substance operating in conscious beings (Papineau 2002; Churchland 2002). Every documented mental phenomenon, from sensory perception to abstract reasoning to emotion to selfhood, has been progressively localized to specific neural substrates whose activity is necessary and, to varying degrees, sufficient for the phenomenon. The principle of causal closure has held up across the full reach of contemporary physics. No plausible candidate for a "mind force" has been identified that would interface with neural matter without leaving experimental traces. Dualism is consequently a position held primarily on philosophical rather than empirical grounds, and the empirical balance of evidence weighs heavily on the materialist side.

Other alternatives. *Panpsychism and Russellian monism.* Panpsychism proposes that all matter possesses some basic form of subjective experience. Contemporary panpsychists hold that some elements of conscious experience can be found in primitive form in some fundamental physical particles, sometimes called "protoconsciousness" or

proto-experiential properties (Chalmers 2015; Goff 2017). The view is held by a small minority of professional philosophers and is widely regarded as a fringe position for two reasons. First, panpsychism faces the *combination problem*: how vast multitudes of micro-experiences combine into the unified macro-experience of a sentient being. Second, it commits to fundamentally non-physical mental properties absent any independent empirical motivation. Physicalists view this as an ontological extravagance that does not earn its keep.

Integrated Information Theory (IIT). IIT departs sharply from the materialist mainstream: it treats consciousness as fundamental and identifies it with integrated information Φ (a measure of the cause-effect power of a system), with what we ordinarily call physical brains being the way conscious systems appear extrinsically to other conscious observers (Tononi 2004; Tononi and Koch 2015; Oizumi et al. 2014). That is, instead of running from physical objects to conscious experience, causality in IIT runs the other way: consciousness is what exists, and its physical substrate is the extrinsic shadow that a Φ-structure casts in the perceptions of other Φ-structures. IIT is a fringe position for several connected reasons. First and most fundamentally, the metaphysical commitment just described is incompatible with materialism, and the physicalist supermajority documented above accordingly rejects it. The position is closer to neutral monism, panpsychism, or a modernized monadology (Tononi himself draws the Leibniz comparison) than to physicalism (Birch 2024). Second, IIT's central axioms are not widely accepted even within the non-materialist literature. The inference from the phenomenal structure of experience to the mathematical structure of Φ is held to be under-motivated (Bayne 2018). Third, Φ is computationally intractable to compute exactly for systems above trivial size, so the theory cannot be directly applied to real brains and must rely on proxies of contested validity. Fourth, the recent "unfolding argument" (Doerig et al. 2019) contends that IIT and other causal-structure theories cannot in principle be distinguished by empirical evidence from functionally equivalent feed-forward systems, raising a serious falsifiability concern.

B.2 The Neural Basis of Valenced Experience

We review here the leading theories of the neural basis of valenced experience. These theories differ in the neural substrate they establish as sufficient for valenced experience. However, they all link valenced experience to brain architecture, agreeing on the implication that similar brain structures support similar valenced experience.

Methodological considerations. Theories of valenced experience are claims about which physical brain mechanisms accompany felt states. The methodology is therefore organized around *markers* that are anatomically independent of the proposed substrate: (i) behavioral markers (conditioned place avoidance, motivational trade-offs, analgesic self-administration, judgment-bias paradigms), (ii) physiological markers (autonomic and HPA-axis responses), and (iii) pharmacological markers (sensitivity to analgesics, anxiolytics, and naloxone). These markers are validated as proxies by their correlation with self-report *in humans*, then carried over by analogy to non-human species. Note that the apparatus is defeasible: no third-person marker is logically sufficient for ascribing a first-person state, an instance of the hard problem of consciousness. But it is the only route available, since the strongest evidence type (verbal report) is unavailable in precisely the cases that bear dispute.

Anatomical primer. Two broad categories of brain structure are candidates for the substrate of valenced experience. On the one hand, evolutionarily ancient *subcortical* structures, in particular the periaqueductal gray (PAG) and the mesolimbic dopamine system. On the other hand, *integrative brain regions*: the cortex in mammals, and functional analogs in other taxa.

Subcortical and cortical processes. It is useful to start with a taxonomy of layers of experience (Tulving 1985). *Anoetic* consciousness refers to raw phenomenal feel, without conceptual or self-referential content. It is sometimes called the basic "what-it-is-likeness" of experience. *Noetic* consciousness refers to experience accompanied by conceptual or propositional content and the recognition that one is experiencing, and *autonoetic* consciousness refers to self-knowing experience embedded in an auto-biographical self that locates the experience in subjective time. The consensus is that anoetic affect is fully realized in subcortical machinery, and the cortex contributes to higher (noetic and autonoetic) layers of experience. The key point of disagreement is whether subcortical-based anoetic consciousness is indeed valenced experience, or whether it needs cortical re-representation to be truly experienced.

One family of theories (Panksepp 1998; Merker 2007; Damasio and Carvalho 2013) holds anoetic consciousness to be valenced experience, and hence that subcortical structures alone suffice for valenced experience.[25] In this view, cortex *elaborates* affect (categorization, semantic content, self-referential framing) and *modulates* it (regulation, voluntary control, re-representation), but does not *generate* the raw affective state.

[25] Anoetic differs from unconscious, e.g., a spinal withdrawal reflex, whose neural mechanism consists of local spinal circuits with no global integration.

The most basic level of valenced feeling is fully realized in subcortical machinery.[26] The competing theories hold that anoetic experience is not experience. In this view, subcortical activity is unconscious processing that needs cortical re-representation to become an experience of any kind. The family includes, among others, the Global Neuronal Workspace theory (Dehaene et al. 2017), higher-order theory (LeDoux and Brown 2017; Rosenthal 2005), and recurrent processing theory (Lamme 2006). All locate consciousness in different brain locations and mechanisms but share the requirement of cortical involvement. Holders of this view acknowledge all the behavioral and physiological evidence for subcortical affective processing, but interpret it as evidence of unconscious processing that needs cortical re-representation to become an experience of any kind.

The disagreement is hard to settle exactly because it instantiates the hard problem of consciousness. Appendix B.3 discusses the evidence in favor of both views. We do not attempt to adjudicate this dispute here.

B.3 Evidence on the Neural Basis of Valenced Experience

Evidence in favor of subcortical sufficiency. Three lines of evidence are taken to support this position. (i) *Deep brain stimulation in humans.* Stimulation of subcortical structures during clinical procedures elicits intense and vivid affective reports (panic, dread, pleasure, sexual feelings) at intensities that cortical stimulation rarely produces (Heath 1972; Mobbs et al. 2007; Panksepp 2005). (ii) *Cross-species behavioral homology.* The same stimulation in homologous midbrain regions of non-human mammals produces behaviors closely matched to the human affective reports (Panksepp 2011). (iii) *Decorticate and decerebrate individuals.* Animals whose cortex has been surgically removed, or whose cortex never developed in early developmental preparations, retain emotional behaviors: distress vocalization, fearful withdrawal, appetitive engagement, play (Panksepp 1998). The strongest empirical version of this argument concerns human hydranencephaly: Merker (2007) documents that children born without a cerebral cortex nevertheless smile, laugh, cry, recognize familiar people, prefer specific songs and toys, and engage in goal-directed behavior at the level appropriate to their motor capabilities, which Merker reads as evidence that the preserved subcortical structures

[26] Some researchers go further. Solms centers consciousness in the brainstem ("consciousness is affect"), which provides the felt quality, while the cortex provides the conceptual and semantic content. Damasio and Carvalho (2013) argues that subcortical mechanisms generate core consciousness (anoetic plus basic body-awareness, here-and-now) while the cortex generates extended consciousness (autobiographical).

suffice for some form of conscious affective experience.

Evidence in favor of cortex-based theories. Three lines of evidence are taken to support the position that conscious valenced experience requires cortical (or cortex-analog) processing. (i) *Subliminal subcortical processing.* Masked fearful faces, presented too briefly for conscious perception, produce amygdala activation comparable to that produced by consciously perceived fearful faces (Whalen et al. 1998), evidence that subcortical affective processing routinely occurs without conscious accompaniment. The same pattern holds for visual perception more broadly: consciously perceived stimuli produce a late, all-or-nothing fronto-parietal "ignition" that subliminal stimuli do not, across attentional-blink and visual-masking paradigms (Sergent et al. 2005; Del Cul et al. 2007; Dehaene 2014). (ii) *Dissociations of conscious feeling from subcortical defensive response.* Feinstein et al. (2011) report that the bilateral-amygdala-lesion patient SM, who reliably fails to show autonomic or behavioral fear responses to standard threat stimuli, nevertheless reports panic and intense fear when subjected to CO_2 inhalation. LeDoux and Brown (2017) read this as direct evidence that the conscious feeling of fear is generated by cortical machinery the amygdala lesion leaves intact, while the amygdala itself produces a defensive survival response without phenomenal accompaniment. (iii) *Cortical disruption tracks loss of consciousness.* Across wake, deep slow-wave sleep, anesthesia, and disorders of consciousness, the perturbational complexity index of Casali et al. (2013), a TMS-EEG measure of cortical integration, reliably discriminates conscious from unconscious states, while subcortical activity is comparatively preserved (Mashour et al. 2020). Vegetative-state patients with preserved cortical activity perform mental-imagery tasks indistinguishably from controls (Owen et al. 2006), suggesting that conscious access tracks cortical preservation even when behavior is absent. Across these three lines, the cortex-required reading is that subcortical activity is sophisticated but unconscious, and that conscious valenced experience requires the cortical or cortex-analog processing emphasized by the Global Neuronal Workspace (Dehaene et al. 2017), higher-order theory (LeDoux and Brown 2017; Rosenthal 2005), recurrent processing theory (Lamme 2006), and the attention schema (Graziano 2013).

The evolutionary roots of valenced experience. The brain-characteristics lens connects welfare range to neural similarity. A second, complementary lens grounds valenced experience in selection pressures. The same selection pressures shape similar brains and similar experiences. Both lenses point in the same direction.

The leading theories of valenced experience treat valence as an evolved adaptation

that serves to motivate fitness-improving actions and discourage fitness-reducing ones. Cabanac (1992) influentially proposed that pleasure functions as a "common currency" in the brain, converting competing motivational claims (hunger, thirst, fear, social bonds) into a single comparable quantity that the organism uses to choose among them. Valenced experience is thought to play three roles in support of this function: it represents fitness-relevant information, it serves as a common currency for decision-making across competing motivations, and it facilitates learning by tagging outcomes as good or bad. The notion of a self evolved in conjunction with the capacity for action: once an animal moves, it becomes useful for it to distinguish stimuli arising from the external world from stimuli generated by its own actions, and the rudiments of an experiencing self become adaptive: a locus that the world impinges on, that initiates action, and that registers the consequences.

If valenced experience is an evolved adaptation, then animals facing similar challenges should have evolved similar valenced responses, supported by similar nervous systems.

B.4 Arguments on Cognition Scaling More Steeply with Neural Architecture than Affect

The conservatism claim splits into two restrictions on the Cobb–Douglas parameters: $\alpha_u \geq \alpha_c$ (the affective production function loads more on subcortical neurons than the cognitive production function does), and $\kappa_u \leq \kappa_c$ (whatever shared substrate exists is used by affect with no greater scale elasticity than cognition uses it). The arguments fall into three categories (anatomical, evolutionary, and functional), and each restriction is supported by arguments in each category. We present these in turn.

B.4.1 Affect Relies More on Subcortical Machinery, Cognition More on Cortical Machinery: $\alpha_u \geq \alpha_c$.

Anatomical separation. Cognition is uncontroversially cortex-dependent: working memory, abstract reasoning, language, and flexible planning load on heteromodal cortex, and the prefrontal expansion of the primate lineage is the standard anatomical anchor for the cognitive capacities tested across species. Affect is contested: as discussed in Appendix B.2, theories disagree on whether subcortical machinery is sufficient or whether cortical re-representation is required. In any case, there is no disagreement that the midbrain plays a central role in generating affect: even for cortex-based theories,

subcortical circuits provide the lower-order unconscious inputs that become conscious experiences in cortical circuits (see, e.g., LeDoux and Brown 2017). In addition, some suggest that the cortical contribution to affect, if present, is concentrated in a small *paralimbic* compartment (anterior cingulate, anterior insula, and orbitofrontal cortex) rather than spread across heteromodal cortex that supports cognition (Craig 2009; Mesulam 1998). Therefore, the affective production function plausibly puts more weight on $z_{m,s}$ relative to $z_{p,s}$ than the cognitive production function does, that is, $\alpha_u \geq \alpha_c$.

Phylogenetic priority. The consensus view is that valenced experience emerged temporally prior to flexible cognition (Panksepp 1998; Damasio 1999, 2010; Damasio and Carvalho 2013). The logic is that the original function of any nervous system is interoceptive monitoring, registering body states and signaling them to action-coordinating circuits, and that a primitive valenced signal is what such monitoring produces. Cognition is built on top of this base, not under it: there is no fitness-relevant problem that cognitive elaboration solves prior to having a valenced signal to act on. Anatomically, the subcortical structures that produce the valenced signal (PAG, hypothalamus, amygdala, brainstem affective nuclei, mesolimbic dopamine) are conserved across all vertebrates, while the cortical expansion that supports flexible cognition is much more recent, accelerated in the primate lineage, and most pronounced in humans. Phylogenetic priority reinforces the anatomical separation.

B.4.2 Scaling: $\kappa_u \leq \kappa_c$.

Even granting any cortical contribution to affect that one wants to allow, three classes of evidence (functional, evolutionary, and anatomical) suggest the cortical scaling of affect is shallower than the cortical scaling of cognition.

Evolutionary argument. Selection on valence is for *adequacy*: once the signal is intense enough to motivate fitness-improving actions, additional cortical machinery does not improve outcomes. Emotions are evolved programs with bounded functional roles, designed to coordinate other psychological mechanisms rather than to admit indefinite refinement (Nesse 1990; Tooby and Cosmides 2008; Nesse 2004). Cognition, by contrast, serves competitive, maximizing functions: outsmarting conspecifics, manipulating the environment, and accumulating cultural knowledge (Dunbar 1998; Reader and Laland

2002; Sol 2009; Burkart et al. 2017).[27] Selection truncated affective intensity at the level required to dominate competing motivations and did not reward additional capacity. Cognitive selection in social and ecological niches has no analogous diminishing-return structure. The asymmetric evolutionary pressure suggests asymmetric scaling: animals with large brains supporting advanced cognition do not need to have proportionally larger welfare ranges for sentience to achieve its evolutionary role.

Functional arguments. *Dimensionality of affect vs. cognition.* Consider a population of N neurons in some brain region. At each moment the population's state is the list of N firing rates, a point in an N-dimensional "neural state space." Over time or over a sequence of stimuli, one can characterize how many independent directions the population state actually varies along, its intrinsic dimensionality, as well as the dimension of the *readout*, i.e., the projection of the population state onto the axes along which downstream circuits decode and act. There is a notable affect-vs-cognition asymmetry in readout dimension. For affect, the readout is low-dimensional. The neuroeconomic literature finds that the brain converts heterogeneous motivational inputs (hunger, fear, social bond, prospective gain) into an essentially one-dimensional value signal read out by the action-selection circuit, the brain's "common currency" posited by Cabanac (1992).[28] The low-dimensional readout is not an artifact of measurement: the action-selection function *requires* a scalar comparator, so a high-dimensional affective readout would fail at the function it serves. For cognition, discrimination, composition, and planning require many independent task-relevant variables to be simultaneously available to downstream circuits. A high-dimensional readout is required to represent distinct concepts, percepts, plans, or inferences. This is the classical Fodor–Pylyshyn argument that systematic, productive thought requires a representational system with many independently varying constituents (Fodor and Pylyshyn 1988). Cortical population codes match this high-dimensionality requirement empirically.[29] In turn, the

[27] Dunbar (1998) documents the social-brain correlation across primates (neocortex ratio scales with social group size). Reader and Laland (2002) extend this to ecological cognition, finding that rates of behavioral innovation, social learning, and tool use covary with executive-brain size. Sol (2009)'s cognitive-buffer hypothesis predicts selection for greater brain capacity precisely where the environment rewards flexible problem-solving. Finally, Burkart et al. (2017) synthesize cross-species evidence for a domain-general *g*-like factor in nonhuman animals.

[28] Levy and Glimcher (2012)'s meta-analysis of fMRI studies identifies a small group of specific brain sites (in the ventromedial prefrontal cortex and orbitofrontal cortex) that appear to encode the subjective values of different types of rewards on a neural common scale. Using converging evidence from neurophysiology, imaging, and lesion studies, Padoa-Schioppa (2011) localizes the same scalar representation in the same brain regions.

[29] Rigotti et al. (2013) record from monkey prefrontal cortex during a cognitive task and find that neurons exhibit *nonlinear mixed selectivity* for combinations of task variables, generating high-dimensional

high-dimensional coding capacity scales with neuron counts. Cortical expansion therefore plausibly raises the cognitive readout's effective dimension, and so its scaling, but cannot raise the affective readout's dimension without breaking the function it serves: bounded for affect, open-ended for cognition.

The dimensional prediction is matched in the cross-species comparative literature: the absolute number of cortical neurons, rather than brain mass or encephalization quotient, is the strongest known neural correlate of cognitive performance across mammals. Herculano-Houzel (2017) synthesizes a decade of isotropic-fractionator counts in support of this conclusion, with the cognition–neuron relationship holding across rodents, primates, and afrotherians and being violated by brain-mass or EQ-based predictors (e.g., elephants and cetaceans have larger brains but fewer cortical neurons than humans, and rank accordingly on cognitive tests). Dicke and Roth (2016) converge on the same conclusion from an independent comparative review.

We can note two related points. First, the conscious channel for affect is unitary, while cognition handles parallelization. The unity-of-consciousness constraint (Bayne 2010) holds that the phenomenal field is integrated: at any moment a subject has a single overall conscious state. For affect this constraint is observed behaviorally: at high intensity, affective states do not co-occur (one cannot simultaneously feel intense fear and intense pleasure), and this corresponds to the finding of the scalar neural representation discussed above. For cognition, by contrast, there is strong evidence of parallelization.[30] Adding cortical capacity multiplies the parallel streams cognition can carry. It does not multiply the affective channel, which is constrained to remain unitary by the function it serves. Second, affect and cognition sit at different ends of the speed–precision trade-off.[31] Affective signaling must be fast, while cognition operates under a much longer time budget. Adding cortical capacity does not relax the speed

population responses whose dimensionality predicts task performance and collapses just before behavioral errors. Fusi et al. (2016) synthesize this evidence and argue that mixed selectivity is the substrate of higher cognition, since high dimensionality enables the linear readout of arbitrarily many task-relevant variables from a single population. Stringer et al. (2019) record neurons in mouse primary visual cortex during natural-image viewing and find that effective dimensionality grows with neuron population size. Haxby et al. (2001) establish the corresponding point in human cortex: object-category information in ventral temporal cortex is carried by distributed, overlapping multi-voxel patterns rather than localized to dedicated channels.

[30] The dorsal and ventral visual streams operate in parallel (Goodale and Milner 1992); the seven-network cortical parcellation of Yeo et al. (2011) documents large-scale networks with substantial functional independence.

[31] Bogacz et al. (2006) provide the canonical formalization of the trade-off. Ratcliff and McKoon (2008) review the empirical fit of the diffusion-decision model across two-choice tasks. Heitz (2014) reviews the broader empirical literature and finds that across species and tasks, faster decisions are systematically less precise, with the trade-off frontier well-characterized at both behavioral and neural levels.

constraint affect must satisfy: a brain with more cortical neurons still has to deliver the affective signal before the action window closes. The trade-off thus caps the useful scaling of the affective channel. Cognition, operating on a longer time budget, is not analogously capped.

Homeostatic regulation and saturation of affect. Affective states are actively regulated back toward a baseline, while cognitive content is not. The hedonic-adaptation literature documents that sustained affective states return to a homeostatic set-point regardless of stimulus magnitude (Brickman et al. 1978; Frederick and Loewenstein 1999; Lyubomirsky 2011). This set-point regulation operates physiologically through autonomic and neuroendocrine feedback. Cognition has no equivalent: knowledge accumulates monotonically, learned skills do not regress to a baseline of ignorance, and working memory does not adapt back to a set-point after a strong cognitive episode. Cortical capacity supports a directly accumulative cognitive substrate, whereas the affective substrate is regulated, and additional capacity is partly absorbed by the regulation rather than translated into additional welfare range.

Relatedly, the neurochemical substrate of affect has hard physical limits that the cognitive substrate does not. For instance, the hedonic "liking" reaction is generated by mu-opioid binding in small, anatomically discrete hotspots in the nucleus accumbens shell and ventral pallidum. Mu-opioid receptors have finite binding capacity and downregulate under sustained activation, the standard mechanism behind opioid tolerance (Koob and Volkow 2010). The evolutionary rationale is that performance is maximized at moderate rather than extreme arousal (Yerkes and Dodson 1908; Mendl 1999). The cortical substrate of cognition is organized differently: representations are sparse and distributed across millions of pyramidal neurons, no single neurotransmitter–receptor system acts as a bottleneck, and encoding is dominated by long-term synaptic plasticity, which is essentially open-ended. Scaling cognitive capacity recruits additional neurons and cortical area (Hill et al. 2010; Van Essen et al. 2019) rather than driving a fixed substrate harder. The bounded biochemistry of the hedonic substrate and the open-ended distributed-plasticity architecture of cognition together suggest a sharper ceiling on $\overline{u}_s$ than on c_s at any given level of cortical capacity.

Anatomical arguments. Finally, *within the cortex*, the cognition-relevant heteromodal compartment scales more steeply across species than the affect-relevant paralimbic compartment across humans, non-human primates, and mice (Hill et al. 2010; Van Essen et al. 2019).

C. Data Appendix

C.1 Population Data

We construct two distinct datasets:

1. For the five taxonomic groups (birds, fish, herptiles, insects, mammals), we have a current population level and a percentage decline since 1970, allowing us to back-calculate a 1970 population.
2. For birds and mammals, we also collect species-level current population counts and construct order- and species-level population changes whenever possible, as detailed below.

C.1.1 Population Changes for Wild Vertebrates: The Living Planet Index (LPI)

The Living Planet Database (Zoological Society of London and World Wildlife Fund 2024) contains 29,559 population data series covering 4,844 species. These population series are gathered from a variety of sources such as journals, online databases, and government reports. For each population series, the rate of change from one year to the next is calculated using generalized additive modeling. Annual rates of change in populations of the same species are then aggregated to the species level and then to higher levels (e.g., order or major taxa) using the geometric mean.

We produce population change series at two levels of aggregation. First, we aggregate species at the level of the four major vertebrate taxa (birds, fish, herptiles, mammals). As recommended by the LPI methodology for aggregation at this level, we first take a simple mean to aggregate species at the taxon$\times$realm level (e.g., Neotropical birds) and then use the provided geographic sampling weights to aggregate from the taxon$\times$realm level to the taxon level. Second, we aggregate species at the order level. We produce order-level estimates only for orders supported by more than 100 distinct monitored population series. This yields order-level estimates for 44 orders (18 fish, 11 birds, 10 mammals, 5 herptiles). The aggregation from the species to the order level is unweighted, as recommended by the LPI methodology for finer levels of aggregation.

Table C.1 shows the coverage of the LPI by major taxa. The coverage is best for mammals and birds, with covered shares around 15%, while fish and herptiles have covered shares below 10%.

We now discuss a number of potential issues with the LPI. First, the species covered

TABLE C.1: LPI Species Coverage

TAXA	SPECIES COVERAGE (%)	SPECIES LPI	SPECIES TOTAL
Birds	14.60 (12.20–14.60)	1,365	9,350 (9,350–11,185)
Fish	6.12 (5.68–7.21)	2,118	34,633 (29,371–37,288)
Herptiles	2.93 (2.68–3.11)	574	19,595 (18,466–21,420)
Mammals	15.22 (11.55–15.22)	787	5,171 (5,171–6,815)

Notes: Total species counts are taken from the Living Planet Report 2024 Technical Supplement, IUCN and Our World in Data. The "Species in LPI" column gives the number of distinct species in the public LPI analytical dataset.

in the LPI may not be representative. This would be particularly problematic if the LPI coverage is systematically biased towards species with larger population losses. The 2024 Living Planet Report Technical document shows that there are more data from regions with a longer history of monitoring and publishing data, such as North America and Europe, compared to other continents or remote or tropical areas. In the existing population series, North America and Europe have experienced the smallest population declines compared with the other continents. When aggregating to the taxon level, this sampling bias is addressed by using the realm weights. McRae et al. (2017) tests for over- or under-representation of species classified as threatened by the IUCN in the LPI. While representation in the LPI is not perfectly balanced, the bias does not systematically go in the direction of better coverage of threatened species. For instance, the most threatened fish and amphibian species tend to be under-represented, while the most threatened reptile species are over-represented. For mammals, there is no significant over- or under-representation. Finally, the LPI may be biased towards species that are easier to track. There is, however, no straightforward way to quantify this bias.

Second, the LPI may be sensitive to outliers (Leung et al. 2020; Murali et al. 2022). To alleviate this concern, we reproduce the LPI index when trimming the population-level changes at different percentiles. This analysis is reported in Table C.2. For birds, fishes, and herptiles, the estimated declines are virtually unchanged. For mammals, we obtain a smaller decline in the 5% tail trimming scenario. Using this smaller decline for mammals does not materially affect any of the paper's conclusions about aggregate welfare changes.

Third, the LPI weights all population series equally when aggregating to the species, order, or major taxon level (except for the geographic reweighting). Underlying population series come in diverse units (density, absolute number, abundance index),

rendering weighting by the relative importance of species impossible.

TABLE C.2: LPI Population-Series Trimming Robustness

Panel A: ±1 capping		
TAXON	BASELINE	
Birds	45.37 [20.56, 61.32]	
Fish	77.16 [65.66, 84.80]	
Herptiles	89.95 [72.38, 96.53]	
Mammals	45.86 [9.89, 70.18]	
Panel B: 1% tail trimming		
TAXON	POOLED	TAXON
Birds	40.76 [15.10, 57.52]	45.58 [20.55, 61.40]
Fish	76.95 [65.65, 84.70]	76.05 [64.37, 83.89]
Herptiles	89.46 [71.08, 96.38]	89.23 [70.57, 96.27]
Mammals	42.69 [6.02, 67.89]	44.27 [7.62, 68.77]
Panel C: 2.5% tail trimming		
TAXON	POOLED	TAXON
Birds	44.36 [20.01, 59.96]	43.69 [19.55, 59.18]
Fish	75.92 [64.14, 83.88]	74.68 [62.34, 82.98]
Herptiles	88.47 [68.72, 95.82]	85.40 [63.08, 94.45]
Mammals	44.28 [2.45, 68.68]	22.52 [-14.92, 52.43]
Panel D: 5% tail trimming		
TAXON	POOLED	TAXON
Birds	46.47 [25.34, 60.36]	46.56 [26.11, 60.87]
Fish	74.01 [62.65, 82.08]	71.55 [58.88, 80.38]
Herptiles	82.40 [55.65, 93.12]	84.80 [61.80, 94.05]
Mammals	20.74 [-11.91, 49.05]	21.41 [-11.46, 49.50]

Notes: This table reports percentage decline from 1970 to 2020, with 95% confidence interval bounds in brackets. Interannual $\log_{10}$ changes are capped at -1 and $+1$ in all specifications.

Validation of the LPI population changes against other sources. Taxon-level percentage declines were checked against the published and grey literature (journal articles, monitoring logs, and conservation-organization reports). We only consider large-scale independent monitoring programs, and do not report studies pertaining to a single species or site. The results of this exercise are recorded in Table C.3. Nearly every available source is narrower than a global taxon-level decline, reporting a trend for specific groups of species or regions. Overall, the orders of magnitude of taxon-level declines reported across studies line up well with the LPI estimates.

TABLE C.3: Validation of LPI Taxon-Level Population Changes Against External Sources

Study	Time span	% Change		Notes
		Total	Annual	
Panel A: Birds				
LPI	1970–2020	−45	−1.19	Weighted
LPI	1970–2020	−25	−0.57	Unweighted
Stanton et al. (2018)	1966–2013	−40	−1.08	Aerial insectivores, North America
Sauer et al. (2020)	1968–2018	−42	−1.08	Grassland specialists, North America
Brlík et al. (2021)	1980–2020	−57	−2.09	Common farmland birds, Europe
Lehikoinen et al. (2018)	2002–2014	−10	−0.87	Mountain specialists, Europe
Bowler et al. (2019)	1990–2015	−13	−0.56	Insectivores, Europe
Kitazawa et al. (2022)	1850–2016	−72	−0.76	Ishikari lowland, Hokkaido, Japan
Wotton et al. (2017)	2009–2015	−39	−7.91	Forest specialists, Uganda
State of India's Birds (2020)	2000–2018	−62	−5.23	Forest specialists, India
State of India's Birds (2020)	2000–2018	−59	−4.83	Grassland/shrub specialists, India
State of India's Birds (2020)	2000–2018	−47	−3.47	Wetland specialists, India
Panel B: Fish				
Study	Time span	% Change		Notes
		Total	Annual	
LPI	1970–2020	−77	−2.90	Weighted
LPI	1970–2020	−64	−2.02	Unweighted
LPI	1970–2020	−47	−1.26	Weighted by mesopelagic population
Pacoureau et al. (2021)	1970–2018	−71	−2.55	Oceanic sharks and rays
Chevalier et al. (2023)	2003–2019	−75	−8.30	Biomass, all species, Lower Mekong basin
Christensen et al. (2014)	1910–2010	−66	−1.07	Biomass, marine predatory fish
Panel C: herptiles				
Study	Time span	% Change		Notes
		Total	Annual	
LPI	1970–2020	−90	−4.50	Weighted
LPI	1970–2020	−83	−3.48	Unweighted
Houlahan et al. (2000)	1966–1997	−47	−2.00	Amphibians, mostly North America and Western Europe
Grant et al. (2016)	–	–	−3.79	Habitat occupancy rate, Amphibians, United States
Adams et al. (2013)	2002-2011	−29	−3.7	Habitat occupancy rate, Amphibians, United States
Rovito et al. (2009)	1975–2007	−94	−8.85	Encouter rates, many species of salamanders, Central America and Mexico
Panel D: Mammals				
Study	Time span	% Change		Notes
		Total	Annual	
LPI	1970–2020	−46	−1.22	Weighted
LPI	1970–2020	−53	−1.50	Unweighted
Gallego-Zamorano et al. (2020)	1992–2015	−40	−2.20	Tropical mammals
Henschel et al. (2014)	1970–2005	−59	−2.52	Large mammals, Africa
Medd et al. (2025)	1987–2021	−71	−3.60	Small mammals, North America

Notes: This table compares LPI taxon-level population changes with estimates from external sources; annualized changes are calculated using the compound annual growth rate over each reported time span.

C.1.2 Other Sources for Population Changes for Wild Animals

Population Change for Insects. Beyond the LPI, we use a number of additional sources for population changes. Our estimate of the decline in the number of insects comes from Dirzo et al. (2014). Dirzo et al. (2014) construct an index of change in terrestrial or semi-terrestrial invertebrate abundance using a methodology similar to the LPI. Their estimate is based on 466 populations of 452 unique species across eight taxonomic groups: the family Aphididae and the orders Coleoptera, Diptera, Hemiptera, Hymenoptera, Lepidoptera, Odonata, and Rhabditida. All of these are insect taxa except Rhabditida, an order of nematodes, so we treat this estimate as representative of insects.

They estimate a 45% decline in insect populations from 1970 to 2008, corresponding to a 1.56% yearly decline on average. This corresponds to the lower range of the estimates in the literature: summarizing data from 6 different studies, Wagner et al. (2021) report 1% to 2% annual declines for terrestrial arthropods. Quantifying local declines across two distinct multi-year research programs in Europe, Wildermuth et al. (2026) get 5.1% and 0.5% yearly losses, and Hallmann et al. (2017) reports a 6.1% annual decline in total flying insect biomass in protected areas.

Our analysis requires a population decline from 1970 to 2020. We use a conservative approach and use the same 45% decline, effectively assuming a 0% decline from 2008 to 2020.

Species-level Population Changes for Mammals. Species-level historical and current totals are available only for a few mammal species in Greenspoon et al. (2025). Finally, we use Ceballos and Ehrlich (2002) for mammal range-contraction and extinction flags.

C.1.3 Sources for Current Population Levels of Wild Animals

Species-level population counts exist only for mammals and birds. For mammals, our primary source is Greenspoon et al. (2025) which consolidates marine-mammal catch-reconstruction, unexploited-population estimates, and expert historical and current estimates for terrestrial mammals. Our secondary source is scraped Wikipedia population counts (Wikipedia 2026b), which gather estimates from various scientific sources (with estimates systematically linked to the source). Finally, we cross-check both of these sources against Greenspoon et al. (2023). For birds, our primary source is the scraped Wikipedia population counts (Wikipedia 2026a). We obtain order-level population counts using a bottom-up approach, i.e., we sum species-level counts at the

order level.

At the level of the broad taxa, Bar-On et al. (2018) supplies population totals for birds, fish (mesopelagic and non-mesopelagic), and arthropods. This source does not provide a count for herptiles, so we rely on the estimate from Tomasik (2018). For mammals, it only provides a biomass estimate, not a population count, so we aggregate our species-level mammal counts to obtain the mammal-level estimate. We obtain a separate population count for insects from Smithsonian Institution (1996) at 1×10^{19}.

For mammals, the species-level population counts are exhaustive (by construction). For birds, the species-level population counts aggregate to 13% of the total from Bar-On et al. (2018).

Validation. We perform a number of validations of the population counts by taxon. For mammals, the sum of the species-level counts from Greenspoon et al. (2025) closely matches that from Wikipedia (2026b) (6.75×10^{10} vs. 5.62×10^{10}). These numbers are also broadly in line with the estimate from Tomasik (2018), based on scaling up density estimates (1×10^{11} to 1×10^{12}). Finally, converting species-level population counts into biomass estimates and aggregating them up, we obtain a biomass of wild mammals at 0.0117 Gt C, very close to the 0.0099 Gt C estimate from Greenspoon et al. (2025). For birds, the group-level total from Bar-On et al. (2018) (3.00×10^{11}) lines up well with the estimate from Callaghan et al. (2021) (5×10^{10}–4.28×10^{11}) based on 9,700 species (approximately 90% of all extant species per their estimate), as well as the estimate of Tomasik (2018) (1×10^{11} to 4×10^{11}). For fish, our estimate is in the same range as Tomasik (2018) (1×10^{13} to 1×10^{15}). For insects, the estimate from Smithsonian Institution (1996) at 1×10^{19} is similar to the estimate of 1.16×10^{19} from Royal Entomological Society (2026). As expected, this is an order of magnitude smaller than the total number of arthropods estimated to be 1.422×10^{20} in Bar-On et al. (2018).

For mammals, we cross-check the species-level population counts across sources. This analysis is reported in Table C.4. We find that the three sources broadly agree, with a median species-level absolute difference around 25% for the three pairwise comparisons.

C.1.4 Combined Dataset: Population Estimates Over Time

For all five broad taxa (mammals, birds, fish, herptiles, and insects), we combine our estimates of population counts with estimates of population changes since 1970 to backcast historical taxon-level populations since 1970.

TABLE C.4: Pairwise Differences in Mammal Population Estimates

SOURCES	% MATCHED	ABSOLUTE % DIFFERENCE	
		MEDIAN	MEAN
Greenspoon et al. (2025) vs. Wikipedia (2026b)	99.34	25.50	25.93
Greenspoon et al. (2025) vs. Wikipedia (2026b)	83.47	26.60	51.18
Greenspoon et al. (2023) vs. Wikipedia (2026b)	85.17	24.84	43.35

Notes: This table compares the mammal population estimates across sources. % Matched is the number of matched species divided by the number of total species in the smaller source. Absolute % difference is calculated for each species as $100\frac{|A-B|}{(A+B)/2}$. We report the mean and median of this statistic across matched species.

For mammals and birds, we additionally construct a more disaggregated dataset: we start from current populations and apply species-level growth rates, order-level growth rates, or growth rates for the broad taxon, depending on data availability. For mammals, the current species-level population data for 4,959 species is exhaustive by construction and this backcasting is done with species-level growth rates for 33 species, order-level growth rates for 4,165, and taxon-level growth rates for the remaining 761 species. For birds, current species-level population data exist for 2,290 species. Birds belonging to these species account for 11.5% of the total bird population. Backcasting is done with order-level growth rates for 1,480 species, and the bird growth rate for the remaining 810 species as well as the bird population not accounted for in our species-level dataset.

C.1.5 Spatial Ranges

Spatial ranges come from two sources. BirdLife International (BirdLife International and Handbook of the Birds of the World 2025) supplies the principal bird ranges, restricted to features flagged extant, native, and resident or breeding (`presence = 1`, `origin = 1`, `seasonal` $\in \{1, 2\}$). IUCN Red List maps (IUCN 2025) supply ranges for mammals, amphibians, reptiles, fish, and insects, together with a small IUCN-only bird supplement.

Both range sources are processed on a common equal-area grid (EPSG:6933) at 25-by-25-km resolution, or 625 km^2 per cell. Empty and non-polygonal features are dropped, and the remaining geometries are cleaned, reprojected, and dissolved by scientific name. Rasterization assigns a cell to a range when the cell center falls inside it (`all_touched = False`), and no fractional polygon–cell overlap is computed.

Populations are assigned by the hierarchy above. A valid species-level `pop_curr` is taken first. Where none exists, the order total is distributed across the unmatched species of that order in proportion to cleaned range area, and where no order total is

available the group total is allocated the same way. This last step carries most of the weight for insects, which enter only as group aggregates and must be spread across the IUCN-covered ranges beneath them. The IUCN-only bird supplement is the one exception: it is restricted to species with a valid species-level population, and receives neither order nor group fallback.

Each species then carries an assigned population and a `pct_decline`, from which the per-cell contribution follows:

$$\begin{aligned} \texttt{density} &= \frac{\texttt{assigned\_population}}{\texttt{total\_cleaned\_range\_area}}, \\ \texttt{density\_change} &= \texttt{density} \times \frac{\texttt{pct\_decline}}{100}, \\ \texttt{count\_change} &= \texttt{density\_change} \times 625, \end{aligned}$$

where positive values of `pct_decline` denote declines and negative values denote increases.

The mapping step reassembles the taxa to match the population data. Amphibians and reptiles are combined into herptiles, and the BirdLife birds are merged with the IUCN-only supplement. Summing the five resulting layers (mammals, birds, fish, herptiles, and insects) yields an all-animals map of absolute population change.

C.1.6 Farmed Animals

We process farmed animals separately from the wild-animal datasets to produce the global, wide-format dataset `pop_domestic_detailed_clean.dta`. It is constructed directly from the 2025 normalized FAOSTAT livestock release (Food and Agriculture Organization of the United Nations 2025). Processing retains records whose element is "Stocks," removes records flagged as missing, and drops the composite items "Cattle and Buffaloes," "Sheep and Goats," and "Poultry Birds" to avoid double counting. We convert values reported in FAOSTAT's "1000 An" unit to individual animals. For bees, FAOSTAT's "No" unit denotes beehives rather than individual insects; this unit is retained and recorded in the variable `pop_unit`. The intermediate area-year file consists of both country- and global-level population estimates.

To construct the final dataset, we retain only FAOSTAT's "World" observations and reshape the data into annual population columns from 1961 through 2024. The resulting dataset contains 18 animal categories (asses, bees, buffalo, camels, cattle, chickens, ducks, geese, goats, horses, mules and hinnies, other birds, other camelids,

FIGURE C.1: Farmed Animal Populations

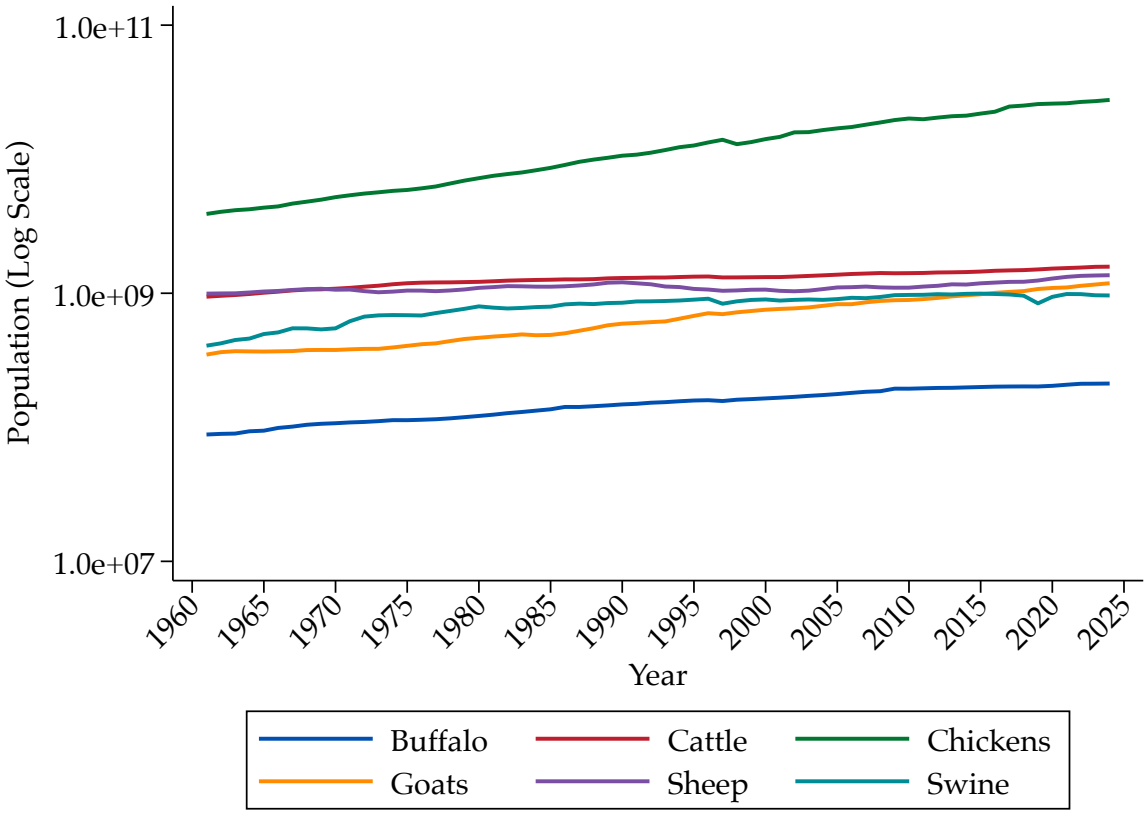

Notes: This figure shows the population of major farmed animals at the global level since 1960. Data are from FAOSTAT.

other rodents, rabbits and hares, sheep, swine, turkeys). Seventeen categories have observations throughout 1961–2024, while the "Other birds" series begins in 1991. Each category is accompanied by a scientific-name field, broad taxonomic group, taxonomic order, and categorical size group, which we assign based on typical weight assumptions: Small (< 10 kg), Medium (10–99 kg), and Large (≥ 100 kg).

Figure C.1 highlights selected farmed animals' global population evolution since 1961.

C.1.7 Additional Results

FIGURE C.2: Wild Animal Population Declines by Taxa

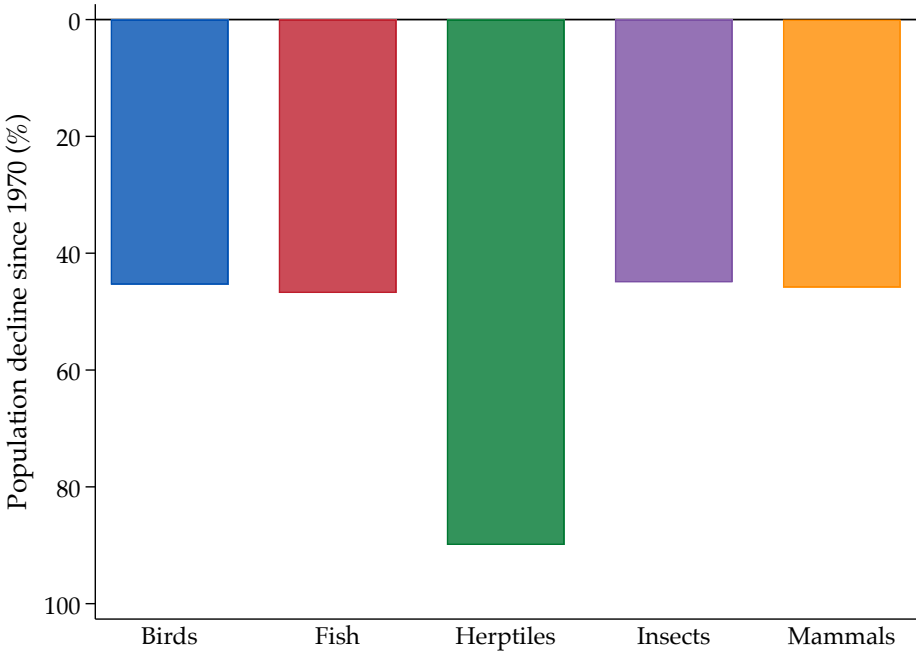

Notes: This figure plots changes in total individuals at the major group level for all wild animals. Each index is calculated using data from the Living Planet Index, except for insects, whose proxy is based on Dirzo et al. (2014).

FIGURE C.3: Global Distribution of Wild-Animal Population Density Change

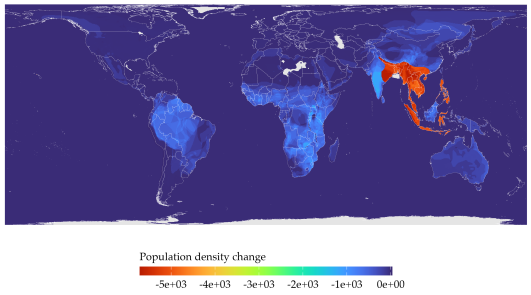

(A) Birds

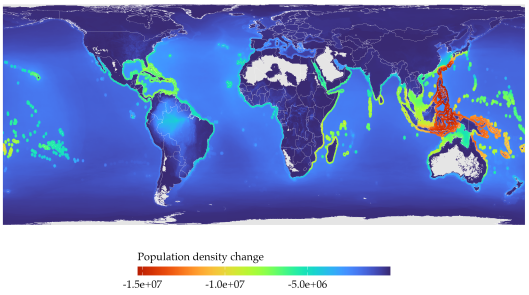

(B) Fish

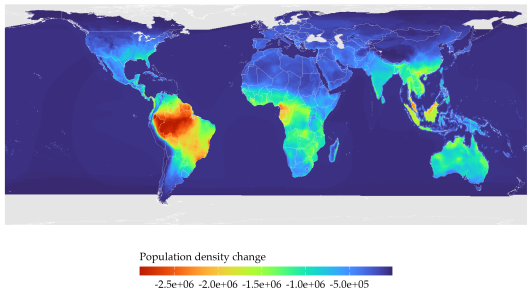

(C) herptiles

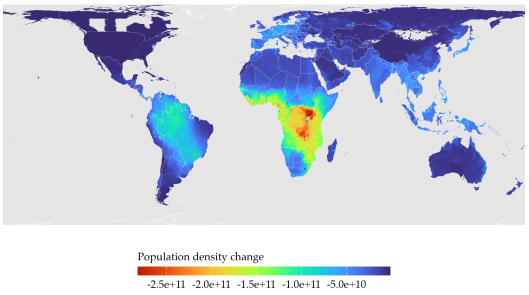

(D) Insects

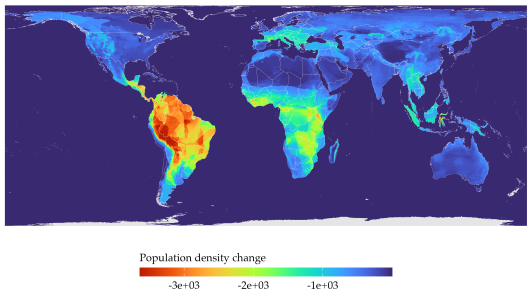

(E) Mammals

Notes: This figure maps unweighted estimated changes in wild-animal population density from 1970 through 2020 for birds, fish, herptiles, insects, and mammals, measured in individuals per square kilometer. Species ranges are from IUCN and BirdLife, and each species' estimated population change is assumed to be spatially uniform within its range. Warmer cells indicate greater estimated density loss. Within each panel, the continuous color scale is linear, with finite nonzero values below the 0.1st percentile and above the 99.9th percentile winsorized to the corresponding boundary colors. Gray cells indicate locations with missing estimates of population-density change.

C.2 Neuron Data

Table C.5 describes the data sources used for neuron counts.

TABLE C.5: Data Sources — Neurons

Source	Description of Dataset
Kverková et al. (2022)	Dataset contains neuron counts for the whole brain, telencephalon, pallium, and other brain regions. Covers bird and reptile species. We use the pallium neuron counts for this dataset.
Wikipedia	Scraped from the Wikipedia page for pallial neurons, only includes pallial neurons but has many species.
Herculano-Houzel et al. (2015)	Contains data on total neurons and cerebral cortex neurons for 39 mammalian species across six clades: Glires, Primates, Scandentia, Eulipotyphlans, Afrotherians, and Artiodactyls.
Sol et al. (2022)	Contains data on total and pallial neurons for 111 bird species.
Olkowicz et al. (2016)	Contains data on total and pallial neurons for 31 bird species.
Marhounová et al. (2019)	Contains data for whole brain and telencephalon neurons in guppies.
Ahrens et al. (2013)	Contains data on total neurons for zebrafish.
Teles et al. (2025)	Contains total and telencephalon neurons for tilapia.
Barron and Mourmourakis (2023)	Contains the number of total neurons for most of the species for which we have cognition tests (A-not-B/cylinder).

FIGURE C.4: Relationship Between Pallial and Total Neurons.

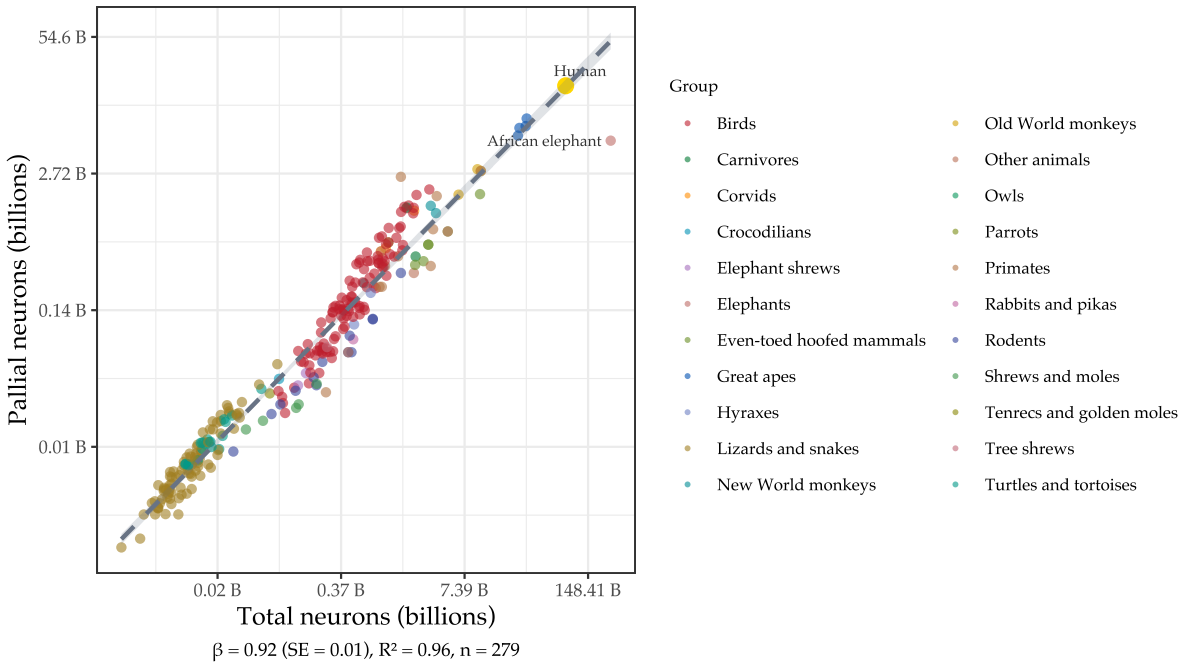

Notes: Relationship between pallial and total neurons for species in which both are observed, on a natural log scale. The dashed line is the fitted line from a regression of log pallial neurons on log total neurons.

TABLE C.6: Assumptions Used in Neurons/Cognition Code.

Assumption short name	**Description**
Manual species coding	Some instances required manual species coding. For example, some Wikipedia entries listed only the genus, as in the honey bee entry, so we manually added a species name. In other cases, the datasets had mistyped names (e.g., *Furoscalops* ↔ *Parascalops*).
Aggregation of dog breeds	The Wikipedia dataset contained pallial neuron counts for multiple dog breeds (the same species, *Canis lupus familiaris*). We take the average of all the dog breeds.
Within species aggregation	More generally, if we have multiple observations for the same species across datasets, we take a within-species average.
Pallial neurons can be predicted from total neurons	We assume that pallial neurons can be predicted from a regression of log pallial neurons on log total neurons. This assumption appears to be reasonable given Figure C.4. For some species, we observe only pallial neurons. For these, we use the inverse mapping.
Total neurons can be predicted from body mass	For some species where we observe neither total nor pallial neurons, we fit a regression of log total neurons on log body mass within each order that has at least three species with observed body mass and total neurons.
Group aggregation	When aggregating to the group level, we assume that the mean or minimum number of neurons that we observe represents the true group mean. This assumption is stronger for groups where we have fewer species and higher populations (e.g. insects).
PCTB human brain neurons	In the PCTB, we assume that a 2.5-year-old human has 80% of adult brain neurons.

C.3 Merged Population and Neuron Data

Mammals. The dataset has one entry per species. The species-level dataset is exhaustive. Each species has:

- A current population count (observed), and a historical population count (obtained by backcasting using the species-level, order-level, or mammal-level growth rate, as described in Appendix C.1.4)
- A neuron count: we use, by order of priority:
  (i) the true or body-mass-predicted neuron count
  (ii) the order mean: species whose order still can't be fitted (too few anchors, or no body mass) inherit the mean neuron count of their order
  (iii) the mammal mean: species in an order with no mean of its own fall back to the mammal class mean

Birds. Birds enter at the order level. The order-level dataset accounts for 11.5% of the bird population. Each bird order has:

- A current population count (obtained by aggregating the observed species counts at the order level), and a historical population count (obtained by backcasting using the order-level, or bird-level growth rate, as described in Appendix C.1.4)
- A neuron count: we use, by order of priority:
  (i) the order mean (the average of the true or body-mass-predicted neuron count for species in the order)
  (ii) the bird mean: species in an order with no mean of its own fall back to the bird class mean

For the birds that cannot be assigned to a given order (88.5% of the total bird population), they are imputed through a single class-wide bird entry carrying the average bird neuron count and decline rate.

Fish, herptiles and Insects. Fish, herptiles and Insects are based on a class-wide average of species neuron counts and a class-wide population decline. These averages yield a count of 2 million neurons for fish, 13 million for herptiles, and 0.2 million for insects.

C.4 Cognitive Tests

TABLE C.7: Data Sources — A-not-B and Cylinder Tests

Source	Description of Dataset
Barron and Mourmourakis (2023)	Aggregates cylinder task scores from MacLean et al. (2014); Kabadayi et al. (2016).
MacLean et al. (2014)	Contains A-not-B task scores.
Lucon-Xiccato et al. (2017)	Contains cylinder task scores for guppies.
Herrmann et al. (2007)	Contains PCTB tests for humans and primates.
Fichtel et al. (2020)	— lemurs.
Schmitt et al. (2012)	— Old World monkeys.
Krasheninnikova et al. (2019)	— parrots.

TABLE C.9: Data Sources — Numeration

Source	Description of Dataset/Paper
Fish	
Lucon-Xiccato et al. (2023)	Study numerosity preferences in newly hatched zebrafish (*Danio rerio*) larvae. They report the preference for larger numerosity for ratios 0.25, 0.33, 0.50.
Bisazza and Santacà (2022)	Study zebrafish (*Danio rerio*) in test of quantities differing by one unit (i.e. 2 v. 3, 3 v. 4, etc.). They show that zebrafish learn up to 4 versus 5 (ratio of 0.8).
Gatto et al. (2020)	Study guppies (*Poecilia reticulata*) trained to choose the larger of two arrays using a Skinner box. Two ratios tested (0.60 = 3 vs 5; 0.75 = 3 vs 4).
Agrillo et al. (2010)	Study mosquitofish (*Gambusia holbrooki*) trained to discriminate large numerosities at ratio 0.50, from 4 versus 8 objects to 100 vs 200. They were able to distinguish between ratios of 0.5 and 0.67, but not 0.75.
Agrillo et al. (2009)	Study mosquitofish trained to discriminate a ratio of 0.67 (2 vs 3).
Stancher et al. (2013)	Study redtail splitfin (*Xenotoca eiseni*) tested across five ratios (0.25, 0.50, 0.50, 0.67, 0.75) in a small-quantity discrimination task.
Insects	

Continued on next page

Table C.9 – continued from previous page

Source	Description of Dataset/Paper
Bortot et al. (2019)	Study honeybees at ratios 2 vs 3 and 3 vs 4. Evidence suggests that bees use absolute rather than relative numerosity.
Howard et al. (2018)	Study honeybees in an ordering setting. They train bees to distinguish "greater-than" and "less-than".
MaBouDi et al. (2020)	Study bumblebees (*Bombus terrestris*) at ratios 0.33 and 0.50.
Bengochea et al. (2023)	Study fruit flies (*Drosophila melanogaster*) at various ratios (0.25-0.75).
Birds	
Rugani et al. (2014)	Study domestic chicks (*Gallus gallus*) at ratios (0.5, 0.67, 0.75).
Rugani et al. (2013)	Study domestic chicks (*Gallus gallus*) at ratio 0.5 (5 vs 10 & 10 vs 20) and 0.67 (6 v. 9)
Tan and Grace (2010)	Study pigeons (*Columba livia*) at ratio 0.33 (2 vs 6, 4 vs 12, 8 vs 24).
Scarf et al. (2011)	Study pigeons in an ordering task (ordered comparisons 1–9).
Ditz and Nieder (2016)	Study crows (*Corvus corone*) in an ordering task.
Pepperberg and Gordon (2005)	Study parrots (*Psittacus erithacus*) in a vocal-labelling task for quantities 1–6, (an ordering-like task).

Continued on next page

Table C.9 – continued from previous page

Source	Description of Dataset/Paper
Reptiles & Amphibians	
Stancher et al. (2014)	Study fire-bellied frogs (*Bombina orientalis*) at ratios 0.50 (1 vs 2) and 0.67 (2 vs 3) in a spontaneous prey-choice paradigm.
Khatiwada and Burmeister (2021)	Study poison dart frogs (*Dendrobates auratus*) at ratio 0.50 (1 vs 2) using live flies as prey.
Gazzola et al. (2018)	Study Hermann's tortoises (*Testudo hermanni*) at ratios 0.25 (1 vs 4) and 0.50 (2 vs 4) in a spontaneous food-choice task.
Lin et al. (2021)	Study Chinese pond turtles (*Mauremys sinensis*) at ratios up to 0.90 (9 vs 10), the hardest pair successfully discriminated.
Mammals	
Tomonaga (2007)	Study chimpanzees (*Pan troglodytes*) at numerosity pairs drawn from 1–8 in a relative numerosity task. Performance follows Weber's Law.
Cantlon and Brannon (2006)	Study rhesus macaques (*Macaca mulatta*) in an ordering task at numerosities 1–9, tested for generalisation to 10–30.
Jordan and Brannon (2006)	Study rhesus macaques (*Macaca mulatta*) at numerosities 1–9 in a delayed match-to-sample task. Performance follows Weber's Law.
Yaman et al. (2012)	Study a bottlenose dolphin (*Tursiops truncatus*) at ratio 0.25 (1 vs 4); fails at the harder ratio 0.75 (3 vs 4).
Miletto Petrazzini et al. (2020)	Study domestic-dog puppies (*Canis lupus familiaris*) at ratio 0.17 (1 vs 6) in a spontaneous food-quantity task.
Gabor and Gerken (2014)	Study Shetland ponies (*Equus caballus*) at ratios 0.50 (1 vs 2) to 0.80 (4 vs 5) in a matching-to-sample task.

Domain	Scale	Task name	Description
Physical	Space	Spatial memory	Locating a reward.
		Object permanence	Tracking of a reward after invisible displacement.
		Rotation	Tracking of a reward after a rotation manipulation.
		Transposition	Tracking of a reward after location changes.
	Quantities	Relative numbers	Discriminating quantity.
		Addition numbers	Discriminating quantity with added quantities.
	Causality	Noise	Causal understanding of produced noise by hidden rewards.
		Shape	Causal understanding of appearance change by hidden rewards.
		Tool use	Using a stick in order to retrieve a reward which is out of reach.
		Tool properties	Understanding of functional and non-functional tool properties.
Social	Social learning	Social learning	Solving a simple but not obvious problem by observing a demonstrated solution.
	Communication	Comprehension	Understanding communicative cues indicating a reward's hidden location.
		Pointing cups	Producing communicative gestures in order to retrieve a hidden reward.
		Attentional state	Choosing communicative gestures considering the attentional state of the recipient.
	Theory of mind	Gaze following	Following an actor's gaze direction to a target.
		Intentions	Understanding what an actor intended to do (unsuccessfully).

TABLE C.8: Description of the Primate Cognition Test Battery Tests.

D. Additional Results: Cognition and Neurons

FIGURE D.1: Cylinder Score and Neurons

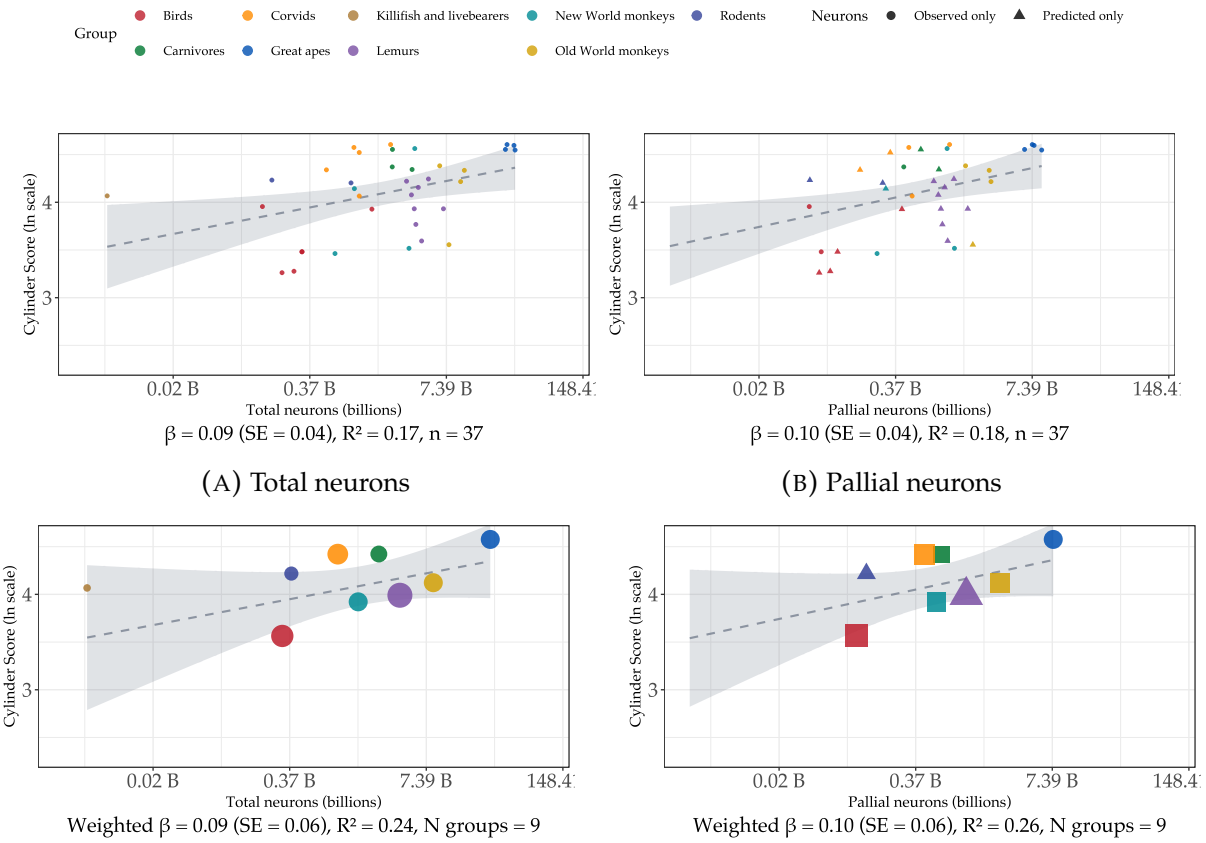

β = 0.09 (SE = 0.04), R^2 = 0.17, n = 37

(A) Total neurons

β = 0.10 (SE = 0.04), R^2 = 0.18, n = 37

(B) Pallial neurons

Weighted β = 0.09 (SE = 0.06), R^2 = 0.24, N groups = 9

(C) Total neurons, binned by group

Weighted β = 0.10 (SE = 0.06), R^2 = 0.26, N groups = 9

(D) Pallial neurons, binned by group

Notes: Cylinder test score versus neuron counts (ln–ln). Panels a, b show individual species. Panels c, d show bins by group, with the binned regression weighted by species count.

FIGURE D.2: A-not-B Score and Neurons

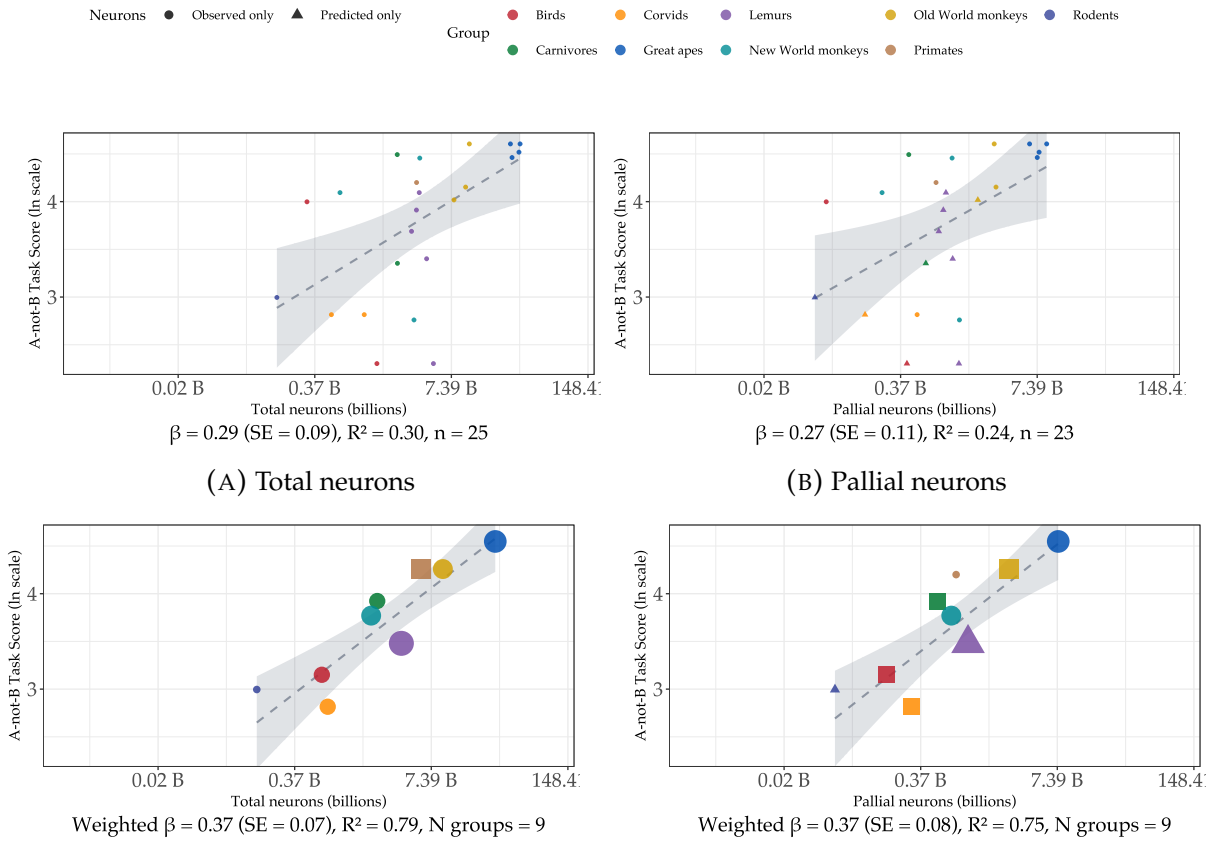

(A) Total neurons

(B) Pallial neurons

(C) Total neurons, binned by group

(D) Pallial neurons, binned by group

Notes: A-not-B test score versus neuron counts (ln–ln). Panels a, b show individual species. Panels c, d show bins by group, with the binned regression weighted by species count.

FIGURE D.3: Numeration Tests

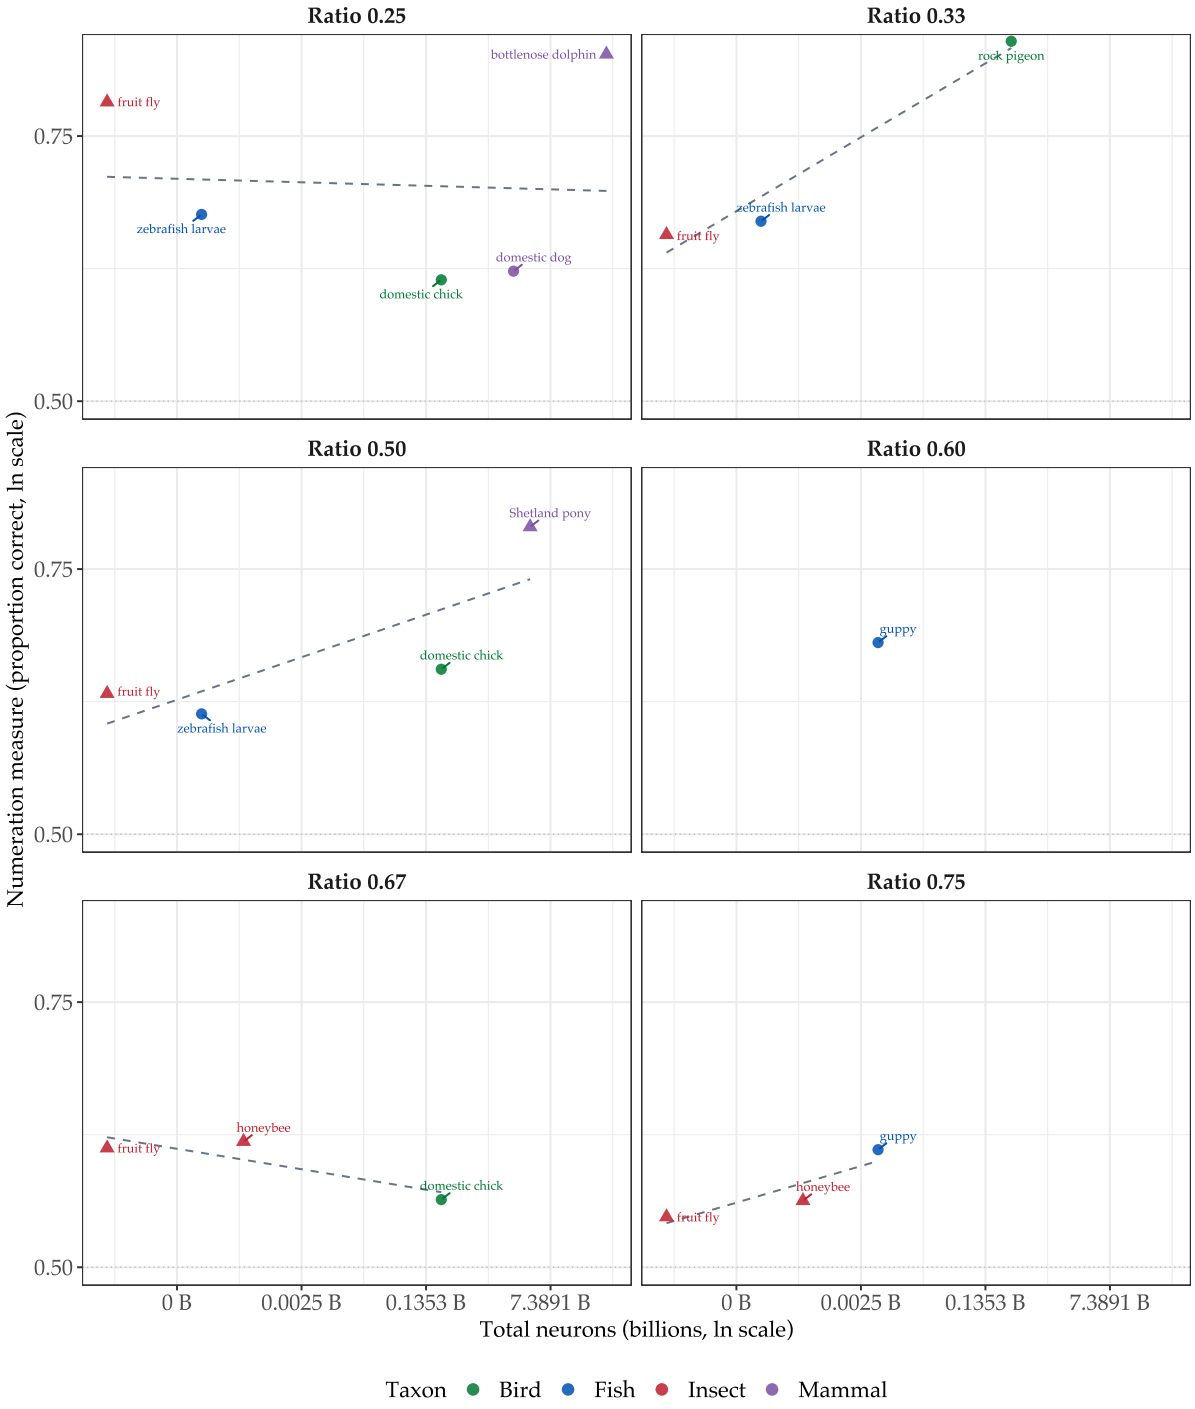

Pooled (across ratios): $\beta = 0.01$ (SE = 0.01), $R^2 = 0.15$, n species = 19

Notes: Relationship between total neurons and test performance on numeration tests at different ratios. The dashed line is the fitted regression line for test performance on log total neurons.

FIGURE D.4: Primate Cognition Test Battery (PCTB)

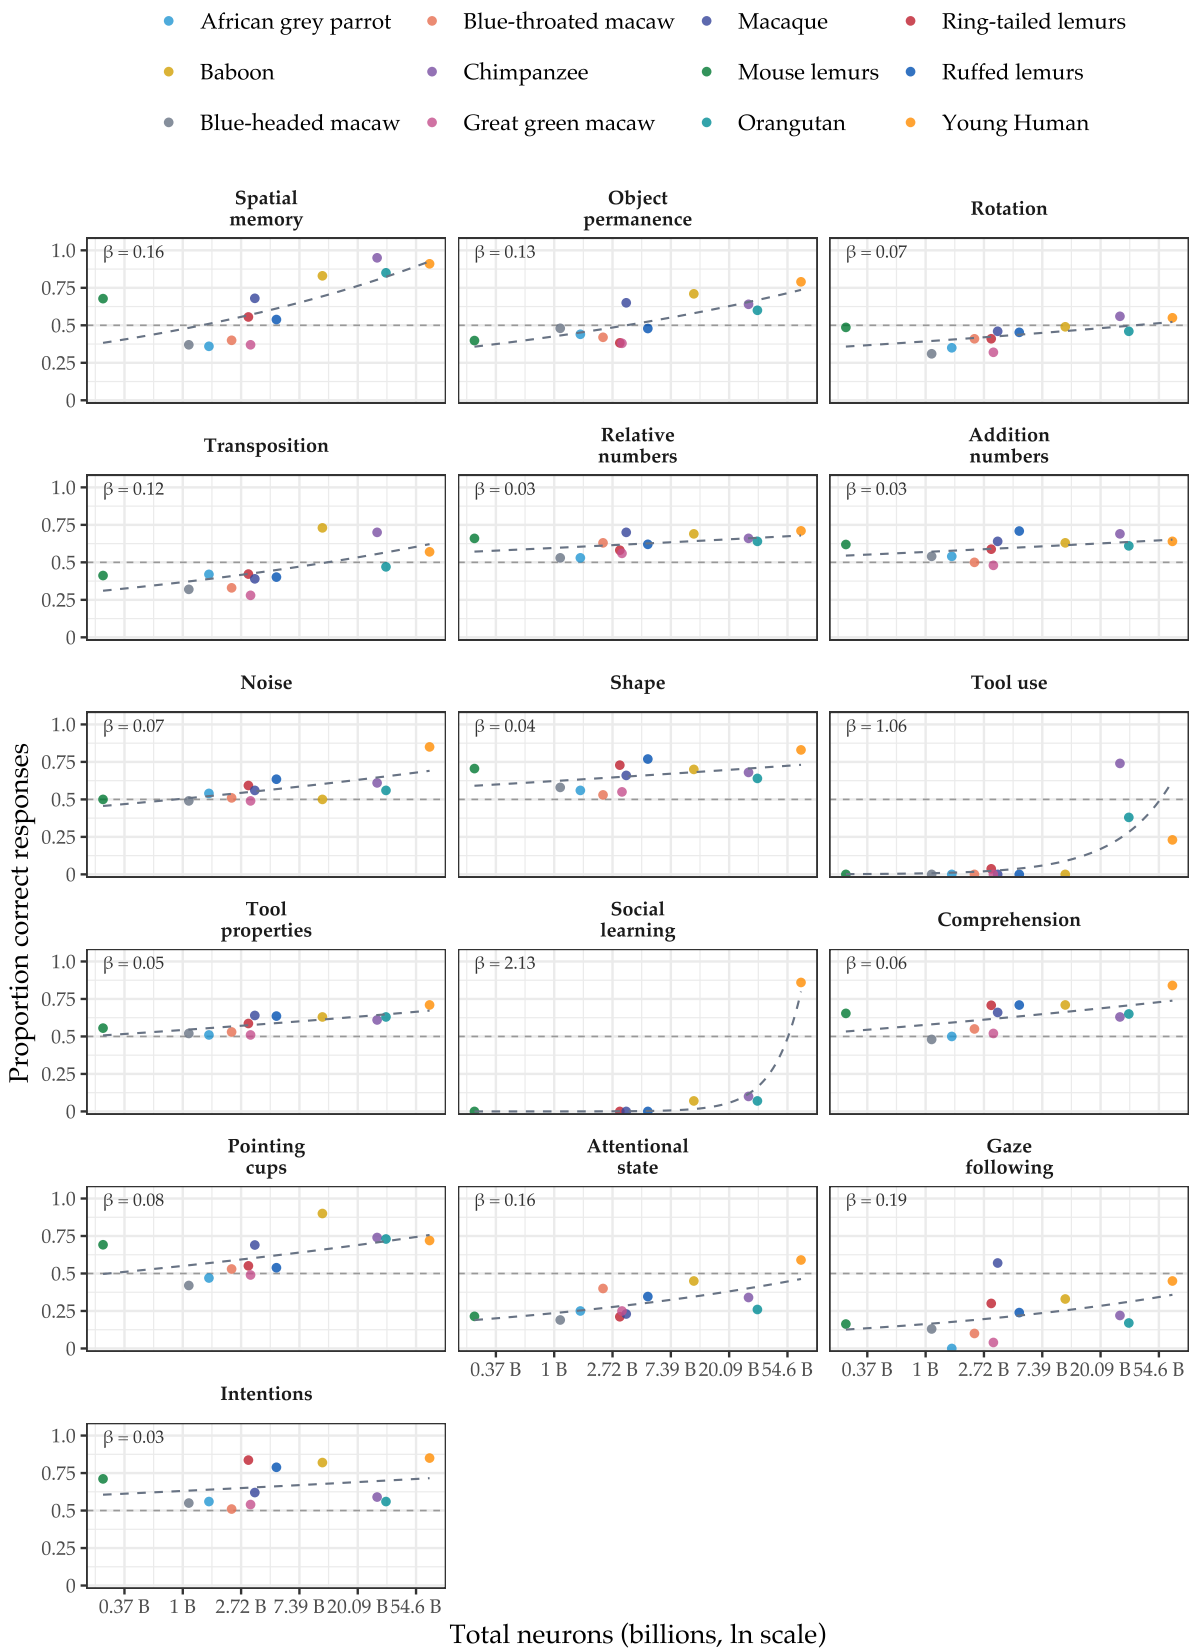

Notes: Relationship between (observed or predicted) total neurons and test performance on each of the PCTB tests. The dashed horizontal line is at 0.5 (passing the test more than 50% of the time). The other dashed line shows the PPML (Poisson Pseudo-Maximum Likelihood) fit of $\mathbb{E}[\text{score}_s|z_s] = \exp(a + \beta \log z_s)$.

E. Additional Results: Welfare

FIGURE E.1: Sentience Losses or Gains by Taxa: Data Aggregated by Taxon

(A) By Group: Mean Neurons by Taxon

(B) Aggregate: Mean Neurons by Taxon

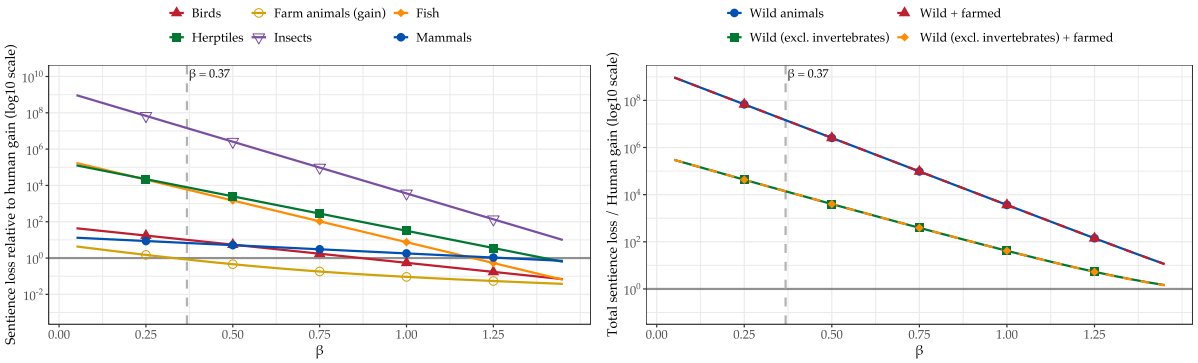

(C) By Group: Minimum Neurons by Taxon

(D) Aggregate: Minimum Neurons by Taxon

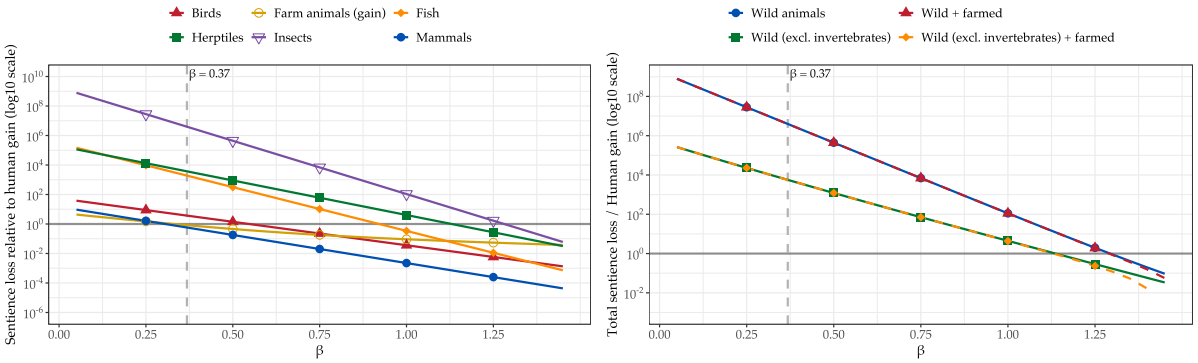

Notes: The y-axis is the cumulative sentience loss or gain for each taxon expressed in aggregate human sentience gains (as defined in (10)). The x-axis is the sentience-to-neurons elasticity β, which determines the relative welfare capacities ω_s. We use relative welfare capacities estimated from total neuron counts. The "Farm animals" category includes all farmed animals, irrespective of taxon. We use population changes by taxon for all taxa. In Panels (A) and (B), we estimate the neuron count for each taxon as the average neuron count among observed species within that taxon. In Panels (C) and (D), we estimate the neuron count for each taxon as the minimum neuron count among observed species within that taxon.

FIGURE E.2: Sentience Losses or Gains by Taxa: Using Pallial Neurons

(A) By Group

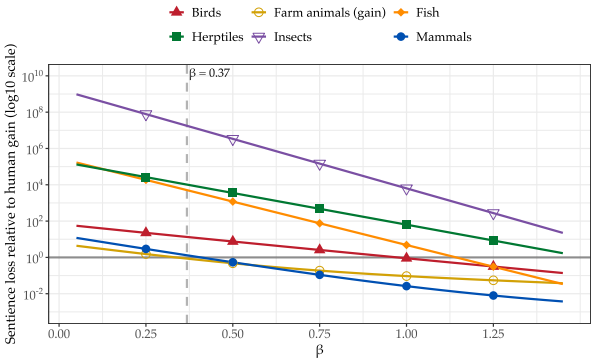

(B) Aggregate

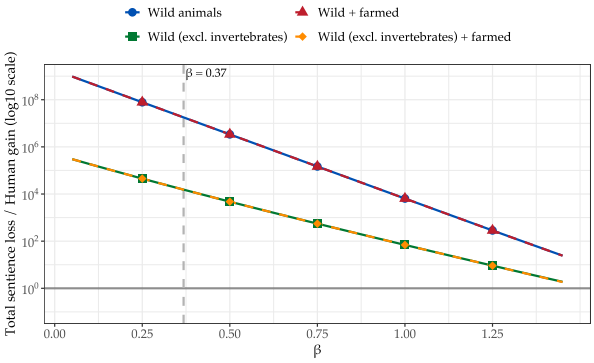

Notes: The y-axis is the cumulative sentience loss or gain for each taxon expressed in aggregate human sentience gains (as defined in (10)). The x-axis is the sentience-to-neurons elasticity β, which determines the relative welfare capacities ω_s. We use relative welfare capacities estimated from pallial neuron counts. The "Farm animals" category includes all farmed animals, irrespective of taxon. We use population data at the most granular level of disaggregation available.

FIGURE E.3: Sentience Losses or Gains by Taxa: Including All Fish

(A) By Group

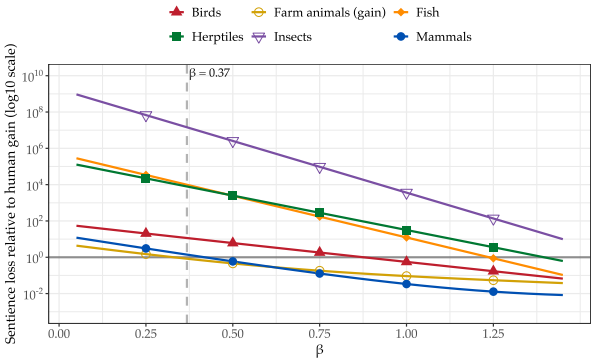

(B) Aggregate

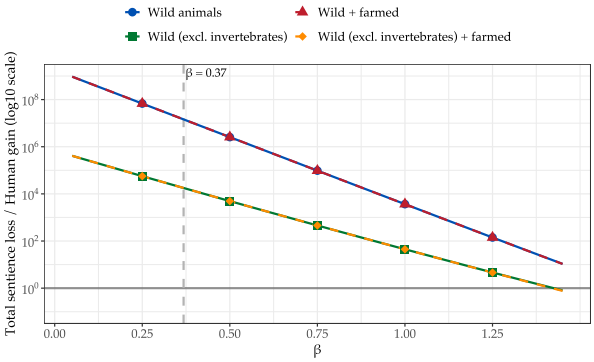

Notes: The y-axis is the cumulative sentience loss or gain for each taxon expressed in aggregate human sentience gains (as defined in (10)). The x-axis is the sentience-to-neurons elasticity β, which determines the relative welfare capacities ω_s. The estimates assume that all fish (non-mesopelagic and mesopelagic) decline in population by the same proportion. The "Farm animals" category includes all farmed animals, irrespective of taxon. We use population data prioritizing the most granular level of disaggregation.

APPENDIX REFERENCES

Adams, Michael J, David AW Miller, Erin Muths, Paul Stephen Corn, Evan H Campbell Grant, Larissa L Bailey, Gary M Fellers, Robert N Fisher, Walter J Sadinski, Hardin Waddle et al., "Trends in amphibian occupancy in the United States," *PloS one*, 2013, *8* (5), e64347.

Agrillo, Christian, Laura Piffer, and Angelo Bisazza, "Large Number Discrimination by Mosquitofish," *PLoS ONE*, December 2010, *5* (12), e15232.

___ **, Marco Dadda, Giovanna Serena, and Angelo Bisazza**, "Use of Number by Fish," *PLoS ONE*, March 2009, *4* (3), e4786.

Ahrens, Misha B., Michael B. Orger, Drew N. Robson, Jennifer M. Li, and Philipp J. Keller, "Whole-brain functional imaging at cellular resolution using light-sheet microscopy," *Nature Methods*, May 2013, *10* (5), 413–420.

Bar-On, Yinon M., Rob Phillips, and Ron Milo, "The biomass distribution on Earth," *Proceedings of the National Academy of Sciences*, 2018, *115* (25), 6506–6511.

Barron, Andrew B. and Faelan Mourmourakis, "The Relationship between Cognition and Brain Size or Neuron Number," *Brain Behavior and Evolution*, July 2023, *99* (2), 109–122.

Bayne, Tim, *The Unity of Consciousness*, Oxford: Oxford University Press, 2010.

___ , "On the Axiomatic Foundations of the Integrated Information Theory of Consciousness," *Neuroscience of Consciousness*, 2018, *2018* (1), niy007.

Bengochea, Mercedes, Jacobo D. Sitt, Veronique Izard, Thomas Preat, Laurent Cohen, and Bassem A. Hassan, "Numerical discrimination in Drosophila melanogaster," *Cell Reports*, 2023, *42* (7), 112772.

Birch, Jonathan, *The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI*, Oxford: Oxford University Press, 2024.

BirdLife International and Handbook of the Birds of the World, "Bird Species Distribution Maps of the World," Version 2025.2 2025. [Online; accessed 2026-06-28].

Bisazza, Angelo and Maria Santacà, "Zebrafish excel in number discrimination under an operant conditioning paradigm," *Animal Cognition*, February 2022, *25* (4), 917–933.

Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D. Cohen, "The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks," *Psychological Review*, 2006, *113* (4), 700–765.

Bortot, Maria, Christian Agrillo, Aurore Avarguès-Weber, Angelo Bisazza, Maria Elena Miletto Petrazzini, and Martin Giurfa, "Honeybees use absolute rather than relative numerosity in number discrimination," *Biology Letters*, 2019, *15* (6).

Bourget, David and David J. Chalmers, "What Do Philosophers Believe?," *Philosophical Studies*, 2014, *170* (3), 465–500.

Bowler, Diana E., Henning Heldbjerg, Anthony D. Fox, Maaike de Jong, and Katrin Böhning-Gaese, "Long-term declines of European insectivorous bird populations and potential causes," *Conservation Biology*, 2019, *33* (5), 1120–1130.

Brickman, Philip, Dan Coates, and Ronnie Janoff-Bulman, "Lottery Winners and Accident Victims: Is Happiness Relative?," *Journal of Personality and Social Psychology*, 1978, *36* (8), 917–927.

Brlík, Vojtěch, Eva Šilarová, Jana Škorpilová, Tomáš Albrecht, Anna Gamero, Alena Kaňuchová, Peter Leopold, Petr Voříšek, and Alena Klvaňová, "Long-term and large-scale multispecies dataset tracking population changes of common European breeding birds," *Scientific Data*, 2021, *8* (1), 81.

Burkart, Judith M., Michele N. Schubiger, and Carel P. van Schaik, "The Evolution of General Intelligence," *Behavioral and Brain Sciences*, 2017, *40*, e195.

Cabanac, Michel, "Pleasure: The Common Currency," *Journal of Theoretical Biology*, 1992, *155* (2), 173–200.

Callaghan, Corey T, Shinichi Nakagawa, and William K Cornwell, "Global abundance estimates for 9,700 bird species," *Proceedings of the National Academy of Sciences*, 2021, *118* (21), e2023170118.

Cantlon, Jessica F. and Elizabeth M. Brannon, "Shared System for Ordering Small and Large Numbers in Monkeys and Humans," *Psychological Science*, 2006, *17* (5), 401–406.

Caplin, Andrew, Mark Dean, Paul W. Glimcher, and Robb B. Rutledge, "Measuring Beliefs and Rewards: A Neuroeconomic Approach," *The Quarterly Journal of Economics*, 08 2010, *125* (3), 923–960.

Casali, Adenauer G., Olivia Gosseries, Mario Rosanova, Melanie Boly, Simone Sarasso, Karina R. Casali, Silvia Casarotto, Marie-Aurélie Bruno, Steven Laureys, Giulio Tononi, and Marcello Massimini, "A Theoretically Based Index of Consciousness Independent of Sensory Processing and Behavior," *Science Translational Medicine*, 2013, *5* (198), 198ra105.

Ceballos, Gerardo and Paul R. Ehrlich, "Mammal Population Losses and the Extinction Crisis," *Science*, 2002, *296* (5569), 904–907.

Chalmers, David J., "Panpsychism and Panprotopsychism," in Torin Alter and Yujin Nagasawa, eds., *Consciousness in the Physical World: Perspectives on Russellian Monism*, New York: Oxford University Press, 2015, pp. 246–276.

Chevalier, Mathieu, Peng Bun Ngor, Kakada Pin, Bunthang Touch, Sovan Lek, Gael Grenouillet, and Zeb Hogan, "Long-term data show alarming decline of majority of fish species in a Lower Mekong basin fishery," *Science of the Total Environment*, 2023, *891*, 164624.

Christensen, Villy, Marta Coll, Chiara Piroddi, Jeroen Steenbeek, Joe Buszowski, and Daniel Pauly, "A Century of Fish Biomass Decline in the Ocean," *Marine Ecology Progress Series*, 2014, *512*, 155–166.

Churchland, Patricia S., *Brain-Wise: Studies in Neurophilosophy*, Cambridge, MA: MIT Press, 2002.

Craig, A. D., "How Do You Feel—Now? The Anterior Insula and Human Awareness," *Nature Reviews Neuroscience*, 2009, *10* (1), 59–70.

Crisp, Roger, *Reasons and the Good*, Oxford: Clarendon Press, 2006.

Cul, Antoine Del, Sylvain Baillet, and Stanislas Dehaene, "Brain Dynamics Underlying the Nonlinear Threshold for Access to Consciousness," *PLoS Biology*, 2007, *5* (10), e260.

Damasio, Antonio, *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*, New York: Harcourt, 1999.

___ , *Self Comes to Mind: Constructing the Conscious Brain*, New York: Pantheon, 2010.

___ **and Gil B. Carvalho**, "The Nature of Feelings: Evolutionary and Neurobiological Origins," *Nature Reviews Neuroscience*, 2013, *14* (2), 143–152.

Dehaene, Stanislas, *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*, New York: Viking, 2014.

___ **, Hakwan Lau, and Sid Kouider**, "What Is Consciousness, and Could Machines Have It?," *Science*, 2017, *358* (6362), 486–492.

Dicke, Ursula and Gerhard Roth, "Neuronal Factors Determining High Intelligence," *Philosophical Transactions of the Royal Society B*, 2016, *371* (1685), 20150180.

Dirzo, Rodolfo, Hillary S. Young, Mauro Galetti, Gerardo Ceballos, Nick J. B. Isaac, and Ben Collen, "Defaunation in the Anthropocene," *Science*, 2014, *345* (6195), 401–406.

Ditz, Helen M. and Andreas Nieder, "Numerosity representations in crows obey the Weber–Fechner law," *Proceedings of the Royal Society B: Biological Sciences*, 2016, *283* (1827), 20160083.

Doerig, Adrien, Aaron Schurger, Kathryn Hess, and Michael H. Herzog, "The Unfolding Argument: Why IIT and Other Causal Structure Theories Cannot Explain Consciousness," *Consciousness and Cognition*, 2019, *72*, 49–59.

Dunbar, Robin I. M., "The Social Brain Hypothesis," *Evolutionary Anthropology*, 1998, *6* (5), 178–190.

Essen, David C. Van, Chad J. Donahue, Timothy S. Coalson, Henry Kennedy, Takuya Hayashi, and Matthew F. Glasser, "Cerebral Cortical Folding, Parcellation, and Connectivity in Humans, Nonhuman Primates, and Mice," *Proceedings of the National Academy of Sciences*, 2019, *116* (52), 26173–26180.

Feinstein, Justin S., Ralph Adolphs, Antonio Damasio, and Daniel Tranel, "The Human Amygdala and the Induction and Experience of Fear," *Current Biology*, 2011, *21* (1), 34–38.

Fichtel, Claudia, Klara Dinter, and Peter M. Kappeler, "The lemur baseline: how lemurs compare to monkeys and apes in the Primate Cognition Test Battery," *PeerJ*, September 2020, *8*, e10025.

Fodor, Jerry A. and Zenon W. Pylyshyn, "Connectionism and Cognitive Architecture: A Critical Analysis," *Cognition*, 1988, *28* (1–2), 3–71.

Food and Agriculture Organization of the United Nations, "FAOSTAT: Crops and Livestock Products (Stocks), Normalized Bulk Download," FAO Statistics Division, Rome 2025. [Online; accessed 2025-12-23].

Frederick, Shane and George Loewenstein, "Hedonic Adaptation," in Daniel Kahneman, Ed Diener, and Norbert Schwarz, eds., *Well-Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation, 1999, pp. 302–329.

Frey, R. G., *Interests and Rights: The Case Against Animals*, Oxford: Clarendon Press, 1980.

Fusi, Stefano, Earl K. Miller, and Mattia Rigotti, "Why Neurons Mix: High Dimensionality for Higher Cognition," *Current Opinion in Neurobiology*, 2016, *37*, 66–74.

Gabor, Vivian and Martina Gerken, "Shetland ponies (Equus caballus) show quantity discrimination in a matching-to-sample design," *Animal Cognition*, 2014, *17* (6), 1233–1243.

Gallego-Zamorano, Juan, Ana Benítez-López, Luca Santini, Jelle P. Hilbers, Mark A. J. Huijbregts, and Aafke M. Schipper, "Combined effects of land use and hunting on distributions of tropical mammals," *Conservation Biology*, 2020, *34* (5), 1271–1280.

Gatto, Elia, Alberto Testolin, Angelo Bisazza, Marco Zorzi, and Tyrone Lucon-Xiccato, "Poor numerical performance of guppies tested in a Skinner box," *Scientific Reports*, October 2020, *10* (1).

Gazzola, Andrea, Giorgio Vallortigara, and Daniele Pellitteri-Rosa, "Continuous and discrete quantity discrimination in tortoises," *Biology Letters*, 2018, *14* (12), 20180649.

Goff, Philip, *Consciousness and Fundamental Reality*, New York: Oxford University Press, 2017.

Goodale, Melvyn A. and A. David Milner, "Separate Visual Pathways for Perception and Action," *Trends in Neurosciences*, 1992, *15* (1), 20–25.

Grant, Evan H Campbell, David AW Miller, Benedikt R Schmidt, Michael J Adams, Staci M Amburgey, Thierry Chambert, Sam S Cruickshank, Robert N Fisher, David M Green, Blake R Hossack et al., "Quantitative evidence for the effects of multiple drivers on continental-scale amphibian declines," *Scientific reports*, 2016, *6* (1), 25625.

Graziano, Michael S. A., *Consciousness and the Social Brain*, Oxford: Oxford University Press, 2013.

Greenspoon, Lior, Eyal Krieger, Ron Sender, Yuval Rosenberg, Yinon M. Bar-On, Uri Moran, Tomer Antman, Shai Meiri, Uri Roll, Elad Noor, and Ron Milo, "The global biomass of wild mammals," *Proceedings of the National Academy of Sciences*, 2023, *120* (10), e2204892120.

___ **, Noam Ramot, Uri Moran, Uri Roll, Rob Phillips, Elad Noor, and Ron Milo**, "The global biomass of mammals since 1850," *Nature Communications*, 2025, *16*, 8338.

Hallmann, Caspar A, Martin Sorg, Eelke Jongejans, Henk Siepel, Nick Hofland, Heinz Schwan, Werner Stenmans, Andreas Müller, Hubert Sumser, Thomas Hörren et al., "More

than 75 percent decline over 27 years in total flying insect biomass in protected areas," *PloS one*, 2017, *12* (10), e0185809.

Haxby, James V., M. Ida Gobbini, Maura L. Furey, Alumit Ishai, Jennifer L. Schouten, and Pietro Pietrini, "Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex," *Science*, 2001, *293* (5539), 2425–2430.

Heath, Robert G., "Pleasure and Brain Activity in Man: Deep and Surface Electroencephalograms during Orgasm," *Journal of Nervous and Mental Disease*, 1972, *154* (1), 3–18.

Heitz, Richard P., "The Speed-Accuracy Tradeoff: History, Physiology, Methodology, and Behavior," *Frontiers in Neuroscience*, 2014, *8*, 150.

Henschel, Philipp, Lauren Coad, Cole Burton, Beatrice Chataigner, Andrew Dunn, David MacDonald, Yohanna Saidu, and Luke T. B. Hunter, "The Lion in West Africa Is Critically Endangered," *PLOS ONE*, 2014, *9* (1), 1–11.

Herculano-Houzel, Suzana, "Numbers of neurons as biological correlates of cognitive capability," *Current Opinion in Behavioral Sciences*, August 2017, *16*, 1–7.

___ **, Kenneth Catania, Paul R. Manger, and Jon H. Kaas**, "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*, September 2015, *86* (3-4), 145–163.

Herrmann, Esther, Josep Call, María Victoria Hernàndez-Lloreda, Brian Hare, and Michael Tomasello, "Humans Have Evolved Specialized Skills of Social Cognition: The Cultural Intelligence Hypothesis," *Science*, September 2007, *317* (5843), 1360–1366.

Hill, Jason, Terrie Inder, Jeffrey Neil, Donna Dierker, John Harwell, and David Van Essen, "Similar Patterns of Cortical Expansion during Human Development and Evolution," *Proceedings of the National Academy of Sciences*, 2010, *107* (29), 13135–13140.

Houlahan, Jeff E, C Scott Findlay, Benedikt R Schmidt, Andrea H Meyer, and Sergius L Kuzmin, "Quantitative evidence for global amphibian population declines," *Nature*, 2000, *404* (6779), 752–755.

Howard, Scarlett R., Aurore Avarguès-Weber, Jair E. Garcia, Andrew D. Greentree, and Adrian G. Dyer, "Numerical ordering of zero in honey bees," *Science*, 2018, *360* (6393), 1124–1126.

IUCN, "The IUCN Red List of Threatened Species. Version 2025-2," 2025. Dataset download. ISSN: 2307-8235.

Jordan, Kerry E. and Elizabeth M. Brannon, "Weber's Law influences numerical representations in rhesus macaques (Macaca mulatta)," *Animal Cognition*, 2006, *9* (3), 159–172.

Kabadayi, Can, Lucy A. Taylor, Auguste M. P. von Bayern, and Mathias Osvath, "Ravens, New Caledonian crows and jackdaws parallel great apes in motor self-regulation despite

smaller brains," *Royal Society Open Science*, April 2016, *3* (4), 160104.

Khatiwada, Sunil and Sabrina S. Burmeister, "Quantity discrimination in a spontaneous task in a poison frog," *Animal Cognition*, 2021, *25* (1), 27–32.

Kim, Jaegwon, *Physicalism, or Something Near Enough*, Princeton: Princeton University Press, 2005.

Kitazawa, Munehiro, Yuichi Yamaura, Masayuki Senzaki, Masashi Hanioka, Haruka Ohashi, Michio Oguro, Tetsuya Matsui, and Futoshi Nakamura, "Quantifying the impacts of 166 years of land cover change on lowland bird communities," *Proceedings of the Royal Society B: Biological Sciences*, 2022, *289* (1975), 20220338.

Koch, Christof, Marcello Massimini, Melanie Boly, and Giulio Tononi, "Neural Correlates of Consciousness: Progress and Problems," *Nature Reviews Neuroscience*, 2016, *17* (5), 307–321.

Koob, George F. and Nora D. Volkow, "Neurocircuitry of Addiction," *Neuropsychopharmacology*, 2010, *35* (1), 217–238.

Krasheninnikova, Anastasia, Roberta Berardi, Mari-Ann Lind, Laurie O'Neill, and Auguste M.P. von Bayern, "Primate cognition test battery in parrots," *Behaviour*, 2019, *156* (5/8), 721–761.

Kverková, Kristina, Lucie Marhounová, Alexandra Polonyiová, Martin Kocourek, Yicheng Zhang, Seweryn Olkowicz, Barbora Straková, Zuzana Pavelková, Roman Vodička, Daniel Frynta, and Pavel Němec, "The evolution of brain neuron numbers in amniotes," *Proceedings of the National Academy of Sciences*, March 2022, *119* (11), e2121624119.

Lamme, Victor A. F., "Towards a True Neural Stance on Consciousness," *Trends in Cognitive Sciences*, 2006, *10* (11), 494–501.

LeDoux, Joseph E. and Richard Brown, "A Higher-Order Theory of Emotional Consciousness," *Proceedings of the National Academy of Sciences*, 2017, *114* (10), E2016–E2025.

Lehikoinen, Aleksi, Lluís Brotons, John Calladine, Tommaso Campedelli, Virginia Escandell, Jiří Flousek, Christoph Grueneberg, Fredrik Haas, Sarah Harris, Sergi Herrando, Magne Husby, Frédéric Jiguet, John Kålås, Åke Lindström, Romain Lorrilliere, Blas Molina, Clara Pladevall, Gianpiero Calvi, Thomas Sattler, and Sven Trautmann, "Declining population trends of European mountain birds," *Global Change Biology*, 12 2018, *25*, 577–588.

Leung, Brian, Anna L Hargreaves, Dan A Greenberg, Brian McGill, Maria Dornelas, and Robin Freeman, "Clustered versus catastrophic global vertebrate declines," *Nature*, 2020, *588* (7837), 267–271.

Levy, Dino J. and Paul W. Glimcher, "The Root of All Value: A Neural Common Currency for Choice," *Current Opinion in Neurobiology*, 2012, *22* (6), 1027–1038.

Lin, Feng-Chun, Martin J. Whiting, Ming-Ying Hsieh, Pei-Jen Lee Shaner, and Si-Min Lin, "Superior continuous quantity discrimination in a freshwater turtle," *Frontiers in Zoology*, 2021, *18* (1).

Lucon-Xiccato, Tyrone, Elia Gatto, and Angelo Bisazza, "Fish perform like mammals and birds in inhibitory motor control tasks," *Scientific Reports*, October 2017, *7* (1), 13144.

___ , ___ , Camilla Maria Fontana, and Angelo Bisazza, "Quantity discrimination in newly hatched zebrafish suggests hardwired numerical abilities," *Communications Biology*, March 2023, *6* (1).

Lyubomirsky, Sonja, "Hedonic Adaptation to Positive and Negative Experiences," in Susan Folkman, ed., *The Oxford Handbook of Stress, Health, and Coping*, New York: Oxford University Press, 2011, pp. 200–224.

MaBouDi, HaDi, H Samadi Galpayage Dona, Elia Gatto, Olli J Loukola, Emma Buckley, Panayiotis D Onoufriou, Peter Skorupski, and Lars Chittka, "Bumblebees Use Sequential Scanning of Countable Items in Visual Patterns to Solve Numerosity Tasks," *Integrative and Comparative Biology*, May 2020, *60* (4), 929–942.

MacLean, Evan L., Brian Hare, Charles L. Nunn, Elsa Addessi, Federica Amici, Rindy C. Anderson, Filippo Aureli, Joseph M. Baker, Amanda E. Bania, Allison M. Barnard, Neeltje J. Boogert, Elizabeth M. Brannon, Emily E. Bray, Joel Bray, Lauren J. N. Brent, Judith M. Burkart, Josep Call, Jessica F. Cantlon, Lucy G. Cheke, Nicola S. Clayton, Mikel M. Delgado, Louis J. DiVincenti, Kazuo Fujita, Esther Herrmann, Chihiro Hiramatsu, Lucia F. Jacobs, Kerry E. Jordan, Jennifer R. Laude, Kristin L. Leimgruber, Emily J. E. Messer, Antonio C. de A. Moura, Ljerka Ostojić, Alejandra Picard, Michael L. Platt, Joshua M. Plotnik, Friederike Range, Simon M. Reader, Rachna B. Reddy, Aaron A. Sandel, Laurie R. Santos, Katrin Schumann, Amanda M. Seed, Kendra B. Sewall, Rachael C. Shaw, Katie E. Slocombe, Yanjie Su, Ayaka Takimoto, Jingzhi Tan, Ruoting Tao, Carel P. van Schaik, Zsófia Virányi, Elisabetta Visalberghi, Jordan C. Wade, Arii Watanabe, Jane Widness, Julie K. Young, Thomas R. Zentall, and Yini Zhao, "The evolution of self-control," *Proceedings of the National Academy of Sciences*, May 2014, *111* (20), E2140–E2148.

Marhounová, Lucie, Alexander Kotrschal, Kristina Kverková, Niclas Kolm, and Pavel Němec, "Artificial selection on brain size leads to matching changes in overall number of neurons," *Evolution*, September 2019, *73* (9), 2003–2012.

Mashour, George A., Pieter Roelfsema, Jean-Pierre Changeux, and Stanislas Dehaene, "Conscious Processing and the Global Neuronal Workspace Hypothesis," *Neuron*, 2020, *105* (5), 776–798.

McRae, Louise, Stefanie Deinet, and Robin Freeman, "The diversity-weighted living planet index: controlling for taxonomic bias in a global biodiversity indicator," *PloS one*, 2017, *12* (1), e0169156.

Medd, Alec, Amanda E. Martin, Adam C. Smith, and Lenore Fahrig, "Continental declines in North American small mammal populations," *Biological Conservation*, 2025, *306*, 111109.

Mendl, Michael, "Performing Under Pressure: Stress and Cognitive Function," *Applied Animal*

Behaviour Science, 1999, *65* (3), 221–244.

Merker, Bjorn, "Consciousness without a Cerebral Cortex: A Challenge for Neuroscience and Medicine," *Behavioral and Brain Sciences*, 2007, *30* (1), 63–81.

Mesulam, M.-Marsel, "From Sensation to Cognition," *Brain*, 1998, *121* (6), 1013–1052.

Mobbs, Dean, Predrag Petrovic, Jennifer L. Marchant, Demis Hassabis, Nikolaus Weiskopf, Ben Seymour, Raymond J. Dolan, and Christopher D. Frith, "When Fear Is Near: Threat Imminence Elicits Prefrontal–Periaqueductal Gray Shifts in Humans," *Science*, 2007, *317* (5841), 1079–1083.

Murali, Gopal, Gabriel Henrique de Oliveira Caetano, Goni Barki, Shai Meiri, and Uri Roll, "Emphasizing declining populations in the Living Planet Report," *Nature*, 2022, *601* (7894), E20–E24.

Nesse, Randolph M., "Evolutionary Explanations of Emotions," *Human Nature*, 1990, *1* (3), 261–289.

___ , "Natural Selection and the Elusiveness of Happiness," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 2004, *359* (1449), 1333–1347.

Nussbaum, Martha C., *Justice for Animals: Our Collective Responsibility*, New York: Simon & Schuster, 2023.

Oizumi, Masafumi, Larissa Albantakis, and Giulio Tononi, "From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0," *PLoS Computational Biology*, 2014, *10* (5), e1003588.

Olkowicz, Seweryn, Martin Kocourek, Radek K. Lučan, Michal Porteš, W. Tecumseh Fitch, Suzana Herculano-Houzel, and Pavel Němec, "Birds have primate-like numbers of neurons in the forebrain," *Proceedings of the National Academy of Sciences*, June 2016, *113* (26), 7255–7260.

Owen, Adrian M., Martin R. Coleman, Melanie Boly, Matthew H. Davis, Steven Laureys, and John D. Pickard, "Detecting Awareness in the Vegetative State," *Science*, 2006, *313* (5792), 1402.

Pacoureau, Nathan, Cassie Rigby, Peter Kyne, Richard Sherley, Henning Winker, John Carlson, Sonja Fordham, Rodrigo Barreto, Daniel Fernando, Malcolm Francis, Rima Jabado, Katelyn Herman, Kwang-Ming Liu, Andrea Marshall, Riley Pollom, Evgeny Romanov, Colin Simpfendorfer, Jamie Yin, Holly Kindsvater, and Nicholas Dulvy, "Half a century of global decline in oceanic sharks and rays," *Nature*, 2021, *589*, 567–571.

Padoa-Schioppa, Camillo, "Neurobiology of Economic Choice: A Good-Based Model," *Annual Review of Neuroscience*, 2011, *34*, 333–359.

Panksepp, Jaak, *Affective Neuroscience: The Foundations of Human and Animal Emotions*, New York: Oxford University Press, 1998.

___ , "Affective Consciousness: Core Emotional Feelings in Animals and Humans," *Conscious-*

ness and Cognition, 2005, *14* (1), 30–80.

___ , "Cross-Species Affective Neuroscience Decoding of the Primal Affective Experiences of Humans and Related Animals," *PLOS ONE*, 2011, *6* (9), e21236.

Papineau, David, *Thinking About Consciousness*, Oxford: Oxford University Press, 2002.

Pepperberg, Irene M. and Jesse D. Gordon, "Number Comprehension by a Grey Parrot (Psittacus erithacus), Including a Zero-Like Concept.," *Journal of Comparative Psychology*, 2005, *119* (2), 197–209.

Petrazzini, Maria Elena Miletto, Fabio Mantese, and Emanuela Prato-Previde, "Food quantity discrimination in puppies (Canis lupus familiaris)," *Animal Cognition*, 2020, *23* (4), 703–710.

Ratcliff, Roger and Gail McKoon, "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks," *Neural Computation*, 2008, *20* (4), 873–922.

Reader, Simon M. and Kevin N. Laland, "Social Intelligence, Innovation, and Enhanced Brain Size in Primates," *Proceedings of the National Academy of Sciences*, 2002, *99* (7), 4436–4441.

Rigotti, Mattia, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi, "The Importance of Mixed Selectivity in Complex Cognitive Tasks," *Nature*, 2013, *497* (7451), 585–590.

Rosati, Connie S., "Persons, Perspectives, and Full Information Accounts of the Good," *Ethics*, 1995, *105* (2), 296–325.

Rosenthal, David M., *Consciousness and Mind*, Oxford: Oxford University Press, 2005.

Rovito, Sean M., Gabriela Parra-Olea, Carlos R. Vásquez-Almazán, Theodore J. Papenfuss, and David B. Wake, "Dramatic declines in neotropical salamander populations are an important part of the global amphibian crisis," *Proceedings of the National Academy of Sciences*, 2009, *106* (9), 3231–3236.

Royal Entomological Society, "Facts and Figures," https://www.royensoc.co.uk/understanding-insects/facts-and-figures/ 2026. Accessed: 2026-07-24.

Rugani, Rosa, Giorgio Vallortigara, and Lucia Regolin, "Numerical Abstraction in Young Domestic Chicks (Gallus gallus)," *PLoS ONE*, 2013, *8* (6), e65262.

___ , ___ , and ___ , "From small to large: Numerical discrimination by young domestic chicks (Gallus gallus).," *Journal of Comparative Psychology*, 2014, *128* (2), 163–171.

Sauer, John R., James E. Hines, and William A. Link, "The North American Breeding Bird Survey, Analysis Results 1966–2018," U.S. Geological Survey data release 2020. Sourced from BirdLife 2022 article.

Scarf, Damian, Harlene Hayne, and Michael Colombo, "Pigeons on Par with Primates in Numerical Competence," *Science*, 2011, *334* (6063), 1664–1664.

Schmitt, Vanessa, Birte Pankau, and Julia Fischer, "Old World Monkeys Compare to Apes in the Primate Cognition Test Battery," *PLOS ONE*, April 2012, *7* (4), e32024.

Sergent, Claire, Sylvain Baillet, and Stanislas Dehaene, "Timing of the Brain Events Underly-

ing Access to Consciousness during the Attentional Blink," *Nature Neuroscience*, 2005, *8* (10), 1391–1400.

Sidgwick, Henry, *The Methods of Ethics*, 7th ed., London: Macmillan, 1907.

Smithsonian Institution, "Numbers of Insects (Species and Individuals)," 1996.

Sol, Daniel, "Revisiting the Cognitive Buffer Hypothesis for the Evolution of Large Brains," *Biology Letters*, 2009, *5* (1), 130–133.

___, Seweryn Olkowicz, Ferran Sayol, Martin Kocourek, Yicheng Zhang, Lucie Marhounová, Christin Osadnik, Eva Corssmit, Joan Garcia-Porta, Thomas E. Martin, Louis Lefebvre, and Pavel Němec, "Neuron numbers link innovativeness with both absolute and relative brain size in birds," *Nature Ecology & Evolution*, September 2022, *6* (9), 1381–1389.

Stancher, G., R. Rugani, L. Regolin, and G. Vallortigara, "Numerical discrimination by frogs (Bombina orientalis)," *Animal Cognition*, 2014, *18* (1), 219–229.

Stancher, Gionata, Valeria Anna Sovrano, Davide Potrich, and Giorgio Vallortigara, "Discrimination of small quantities by fish (redtail splitfin, Xenotoca eiseni)," *Animal Cognition*, January 2013, *16* (2), 307–312.

Stanton, R.L., C.A. Morrissey, and R.G. Clark, "Analysis of trends and agricultural drivers of farmland bird declines in North America: A review," *Agriculture, Ecosystems & Environment*, 2018, *254*, 244–254.

State of India's Birds, "State of India's Birds, 2020: Range, Trends and Conservation Status," Report and associated assessment dataset 2020. Sourced from BirdLife 2022 article.

Stringer, Carsen, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D. Harris, "High-Dimensional Geometry of Population Responses in Visual Cortex," *Nature*, 2019, *571* (7765), 361–365.

Tan, L. and R. C. Grace, "Discrimination and representation of relative numerosity in a bisection task by pigeons," *Learning & Behavior*, 2010, *38* (4), 408–417.

Teles, Magda C., Gonçalo M. Melo, Suzana Herculano-Houzel, and Rui F. Oliveira, "Larger Fish Have Larger Brains With More Neurons Across but Not Within Cohorts Raised in Different Growth Conditions," *Journal of Comparative Neurology*, 2025, *533* (9), e70090.

Tomasik, Brian, "How Many Wild Animals Are There?," Essays on Reducing Suffering 2018. [Online; accessed 2026-05-20].

Tomonaga, Masaki, "Relative numerosity discrimination by chimpanzees (Pan troglodytes): evidence for approximate numerical representations," *Animal Cognition*, 2007, *11* (1), 43–57.

Tononi, Giulio, "An Information Integration Theory of Consciousness," *BMC Neuroscience*, 2004, *5*, 42.

___ and Christof Koch, "Consciousness: Here, There and Everywhere?," *Philosophical Transactions of the Royal Society B*, 2015, *370* (1668), 20140167.

Tooby, John and Leda Cosmides, "The Evolutionary Psychology of the Emotions and Their

Relationship to Internal Regulatory Variables," in Michael Lewis, Jeannette M. Haviland-Jones, and Lisa Feldman Barrett, eds., *Handbook of Emotions*, 3rd ed., New York: Guilford Press, 2008, pp. 114–137.

Tulving, Endel, "Memory and Consciousness," *Canadian Psychology*, 1985, *26* (1), 1–12.

Wagner, David L, Eliza M Grames, Matthew L Forister, May R Berenbaum, and David Stopak, "Insect decline in the Anthropocene: Death by a thousand cuts," *Proceedings of the National Academy of Sciences*, 2021, *118* (2), e2023989118.

Whalen, Paul J., Scott L. Rauch, Nancy L. Etcoff, Sean C. McInerney, Michael B. Lee, and Michael A. Jenike, "Masked Presentations of Emotional Facial Expressions Modulate Amygdala Activity without Explicit Knowledge," *Journal of Neuroscience*, 1998, *18* (1), 411–418.

Wikipedia, "Lists of Birds by Population," https://en.wikipedia.org/wiki/List_of_birds_by_population 2026.

___, "Lists of Mammals by Population," https://en.wikipedia.org/wiki/Lists_of_mammals_by_population 2026.

Wildermuth, Benjamin, Maximilian Bröcher, Emma Ladouceur, Sebastian T Meyer, Holger Schielzeth, Michael Staab, Rafael Achury, Nico Blüthgen, Lionel Hertzog, Jes Hines et al., "Arthropod species loss underpins biomass declines," *Nature Ecology & Evolution*, 2026, *10* (1), 83–94.

Wotton, Simon, Mark Eaton, Danaë Sheehan, Fred Munyekenye, I.J. Burfield, S.H.M. Butchart, K. Moleofi, Dianah Nalwanga, Paul Ndang'ang'a, D. Pomeroy, K.J. Senyatso, and Richard Gregory, "Developing biodiversity indicators for African birds," *Oryx*, 11 2017, *54*, 62–73.

Yaman, Sevgi, Annette Kilian, Lorenzo von Fersen, and Onur Güntürkün, "Evidence for a Numerosity Category that is Based on Abstract Qualities of "Few" vs. "Many" in the Bottlenose Dolphin (Tursiops truncatus)," *Frontiers in Psychology*, 2012, *3*.

Yeo, B. T. Thomas, Fenna M. Krienen, Jorge Sepulcre, Mert R. Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L. Roffman, Jordan W. Smoller, Lilla Zöllei, Jonathan R. Polimeni, Bruce Fischl, Hesheng Liu, and Randy L. Buckner, "The Organization of the Human Cerebral Cortex Estimated by Intrinsic Functional Connectivity," *Journal of Neurophysiology*, 2011, *106* (3), 1125–1165.

Yerkes, Robert M. and John D. Dodson, "The Relation of Strength of Stimulus to Rapidity of Habit-Formation," *Journal of Comparative Neurology and Psychology*, 1908, *18* (5), 459–482.

Zoological Society of London and World Wildlife Fund, "Living Planet Index Database," 2024. Accessed: 2026-05-20.