7

QUANTIFYING ANIMAL WELL-BEING AND OVERCOMING THE CHALLENGE OF INTERSPECIES COMPARISONS

Mark Budolfson and Dean Spears

Animals,[1] like humans, experience different levels of well-being depending on decisions made by others. As a result, the well-being of animals must be included in any full accounting of the well-being consequences of decisions. However, this is almost never done in large-scale policy analyses and investment analyses, even though it is common to quantify the consequences for human welfare in these decision analyses. This is partly due to prejudice, but increasingly also because we do not currently have good methods for quantifying animal well-being consequences and putting them on the same scale as quantified human well-being consequences. We might call this 'the problem of interspecies comparisons.' This important barrier to including animal well-being in decision-making is the result of an insufficiently developed theory and practice of animal well-being and its relation to human well-being.

This handbook chapter explains the problem of interspecies comparisons, explains recent research that develops methods to overcome this problem, and includes animal welfare in rigorous policy and investment analysis (e.g., in analyses of optimal public policies, analyses of optimal philanthropic investment, and so on). The development of these methods is important: incorporating animal welfare in decision analyses would have an important impact on estimates of what prosocial investments of time and money should be made by individuals,[2] businesses,[3] and charities (including for purposes of 'effective altruism'),[4] and similarly for estimates of optimal public policies for correcting market failures (where the full cost of goods is not reflected in their market price),[5] for sustainable intensification of agriculture that aims to take animal welfare into account (producing more food while reducing the overall impacts of agriculture),[6] for climate change policy (how quickly we should be reducing greenhouse gas emissions),[7] and for wilderness protection policy and other challenges related to natural resource management.[8] In all these cases, if the well-being of animals is taken more fully into account, then decisions by individuals and governments will become better, on utilitarian grounds, and more compassionate toward the plight of animals.

Anthropocentrism and Current Economic Analysis: Animal Welfare as Valuable Only Insofar as It Is Valued by Humans

Anthropocentrism is the view that what makes outcomes better or worse is ultimately entirely a matter of their consequences for humans.[9] On this view, animal welfare is not ultimately valuable per se; instead, animal welfare is only valuable insofar as it is ultimately valued by humans. From this perspective, which is dominant in economic policy analysis, the important problem related to animal well-being is that the marketplace and existing methods of policymaking tend to ignore much of the instrumental value of animal welfare to humans, and thus do not take animal welfare into proper account even from an anthropocentric perspective.

To see how there is a tendency to ignore even the instrumental value of animals, consider an analogy to majestic stands of old-growth forests: until recent decades, the fate of old-growth forests in many nations was entirely determined by a marketplace that did not take into account human society's preference for preserving the most majestic of old-growth forests because that preference was not reflected in market prices for timber. As a consequence, there was a 'market failure' in which the outcome determined by the marketplace was inferior[10] for humans by their own lights than the outcome that would have obtained if human society's preferences for preservation had been incorporated into the price of the most majestic of old-growth forests. This is a classic example of market failure deriving from the existence of 'negative externalities' (costs to society that are not internalized in market prices), which shows how externalities can cause free market transactions to lead to suboptimal outcomes for human society by its own lights. For the purposes of this chapter, the key thing to note is that, in these cases, an outcome is suboptimal from an anthropocentric perspective because market prices do not always reflect everything that is of value to humans. Market failure is a very real and widespread phenomenon, and does not essentially depend on the question of how to value nature, as market failure can arise anytime parties external to a transaction bear costs not internalized by the price mechanism or the parties to the transaction, even when no aspect of nature is in play. The current point is that sometimes externalities do indeed exist for elements of nature, including animal welfare. That is, there are cases where animal welfare is valued by humans in a way that is not reflected in unregulated market prices. In those cases, market failure threatens, and outcomes may be worse for humans by their own lights than other feasible outcomes.[11]

With that background in hand, we now note that there is empirical evidence that the current marketplace and public policies do not adequately reflect the value that humans assign to animal welfare; thus, something is going wrong even by the lights of an anthropocentric view that maintains that animal welfare is valuable only insofar as it is valuable to humans.[12] The relevant empirical evidence here is analogous to the old-growth-forests example: economic analysis indicates that sufficiently many humans are willing to pay to improve animal welfare above current levels in a way that implies that the outcome could be made better by humans' own lights by properly taking that willingness to pay into account. As an analogy, imagine a situation where a single unimportant factory is billowing noise and filth into the air of our community, and many of us are willing to pay to prevent it from doing so. In such a case, the amount we are willing to pay could be more than enough to make the factory owners happy to reduce emissions dramatically if our payments are transferred to them, and we would be better off if we did so because we'd prefer that outcome to continuing to suffer the pollution. So there is an opportunity to make some humans better off without making anyone worse off—clearly a better outcome by the lights of anthropocentrism. The empirical argument of many economists who study these issues is that the same is true regarding animal welfare: we can make all humans better off and none worse off by improving animal welfare in specific ways.[13] A complementary argument for the same conclusion is that some targeted animal welfare improvements would more than pay for themselves by reducing the expected harm to human health from diseases, antimicrobial resistance, and the like, where these harms to human health are not reflected in the market prices of animal products; thus, policies that included targeted animal welfare improvements could yield benefits for everyone in expectation.[14]

The methods of economic analysis that underlie these conclusions are methods of estimating two different categories of anthropocentric value, namely, the (anthropocentric) *use value* of animals (human willingness to pay to use animals) and their *nonuse value*. Use value includes willingness to pay for direct use of animals as well as indirect use, including *ecosystem services* such as the value of pollinators in human agriculture, the value of aquatic mollusks in cleaning water for human use, the value of wildlife to human recreation, and so on. Nonuse value includes the willingness to pay merely for an outcome that includes the existence of animals without their use by humans (*existence value*), as well as the *option value* of keeping animals around for potential future human use.

Substantive methods are needed to estimate these anthropocentric values of animals in many cases, as their values are often not readily reflected in market prices,[15] especially when they have the properties of public rather than private goods. This is almost always the situation in connection with nonuse value, which is why the nonuse value of animals is generally ignored in the marketplace and in policy analysis. At the same time, there are widely known methods for estimating nonuse value, namely, contingent valuation studies and revealed preference methods. *Contingent valuation studies* are generally surveys that elicit self-reported willingness to pay to avoid or bring about particular outcomes. Based on respondents' answers, willingness to pay for nonuse value is estimated and can then be incorporated in decision analyses. However, there are many objections to this method, perhaps the most important of which is the worry that it leads to biased and inflated estimates of willingness to pay, because people do not have to back their answers with real monetary investments.[16] There is some evidence that this leads to a large upward bias in estimates of willingness to pay in comparison to the actual choices that people make when real money is on the line.

This points toward the main alternative method, estimates based on *revealed preferences* (i.e., the valuation implicit in actual choices), often based on the valuation implicit in existing regulation, such as applications of the Endangered Species Act. Using the values that are implicit in existing regulation mitigates a worry that would arise if an analysis were to use the values implicit in individual expenditures instead: namely, a societal collective action problem could explain low expenditure by individuals. This is because the nature of a problem might be such that it would not be mitigated by individual expenditures; it would only be mitigated by widespread collective investment. In such circumstances, the absence of collective investment would make individual investment irrational even if individuals preferred a substantial level of investment by all.[17]

These economic methods allow estimation of different species' value to humans, and are the most widely used methods for doing so. However, most philosophers and many economists would insist that even if perfectly developed, these methods can never be adequate for valuing animal welfare, since by their nature they are incapable of assigning any fundamental value to the well-being of animals. Because it is clear that animals experience well-being that humans are often unwilling to pay to protect, this is unacceptable from the point of view that well-being is the fundamental value that economic analysis aims to promote. This is the point of departure for the next section, which introduces the challenges of both estimating animal well-being and making comparisons between animals and humans.

Animal Welfare as Fundamentally Valuable, and the Problem of Interspecies Comparisons

Many researchers from across disciplines and many citizens believe that an important problem with even the best anthropocentric methodology is that the valuation of animals and other aspects of nature within such a methodology is always merely valuation in terms of what only *humans* value.[18] In other words, valuation is always in terms of the ultimate value of outcomes to humans only, and thus assigns no fundamental value to the well-being of animals. For example, on even the best anthropocentric approach, the deaths of billions of birds due to climate change would have disvalue only insofar as the deaths of those birds have disvalue to humans. But many would object that this way of valuing animal lives is fundamentally incorrect because it ignores the value of the birds' own well-being irrespective of its contribution to human well-being. (Similarly, anthropocentrism assigns no fundamental value to the health of ecosystems as holistic entities, which is a separate criticism that we set aside here.[19])

So, according to critics of anthropocentrism, what is needed is the inclusion of the fundamental value of animals' own well-being, even if it is not valued by humans. In what follows, we side with the critics of anthropocentrism on this issue. We agree that since animals experience well-being and humans often do not value that well-being, any coherent welfarist approach must acknowledge that there is an important question of how to incorporate this ignored fundamental value of animal welfare into decision analysis. Although we frame the discussion in terms of welfarist consequentialism for ease of exposition, the structure of the approach that follows is also compatible with many deontological approaches.[20]

Although most philosophers and an increasing number of practitioners agree that anthropocentrism should be rejected, they also tend to agree that there have not been good methods for quantifying animal well-being consequences and putting them on the same scale as quantified human well-being consequences in a decision analysis. This is 'the problem of interspecies comparisons.'

Recent work by Kevin Wong (Wong 2016) has clarified the most difficult problem that needs to be solved in connection with interspecies comparisons. As Wong notes, the key problem is how to estimate the *well-being capacity* (well-being potential) of members of a nonhuman species relative to the well-being capacity of humans. If we knew how to make those interspecies comparisons of well-being capacity, then we could integrate animal welfare consequences into existing methods of decision analysis. This integration would be made possible by deriving empirically based estimates of animal welfare consequences on the same scale as human consequences that typically underpin welfarist decision-making analyses.

For example, suppose an additional degree of climate change will cause us to lose one million life-years of a particular species of bird, and we want to value this on the same scale as the losses to humans from an additional degree of warming that are (let's assume) already modeled and valued based on an assumption about the value of one average human life-year. Wong's point is that if we had a good estimate of the well-being capacity of that species of bird relative to a human, we could then multiply that estimate by the purely empirical impact estimate of one million life-years lost. This would yield an estimate of the amount of well-being lost by that bird species on the same value scale as the existing estimate of human well-being loss, assuming that one degree of additional climate change does not change the quality of life of those birds. And if one additional degree of warming does diminish the quality of life of the remaining birds of that species, we can simply multiply the number of remaining bird species life-years by a further *quality-of-life adjustment* term that is itself an empirical impact estimate from zoological experts and the like. (We can also use such a term to take into account any antecedent diminishment in the well-being experienced by all of the birds including those that would die before the warming.)

This line of thought leads to the following equation for estimating on a single scale the average well-being experienced by a member of a species s (which we symbolize as $\overline{u}_s$) as a function of the average well-being capacity per unit of time of members of s relative to humans $(\overline{\pi}_s)$, multiplied by the average duration of a life of a member of s $(\overline{\delta}_s)$, multiplied by a quality of life adjustment term that estimates how well members of *s* are typically flourishing relative to their species capacity (i.e., a quality of life adjustment term) $(\overline{f}_s)$[21]:

1. $\overline{u}_s = \overline{\pi}_s * \overline{\delta}_s * \overline{f}_s$

Wong's contribution is to highlight the term $\overline{\pi}_s$ as the key unknown term (as the other terms $\overline{\delta}_s$ and $\overline{f}$ are susceptible to existing empirical methods),[22] where the unsolved problem of how to estimate $\overline{\pi}_s$ is the essence of the challenge of interspecies comparisons.

Wong uses formalisms such as Equation 1 to estimate the valuation of animal well-being that is implicit in some decisions within effective altruism that allocate scarce resources between opportunities to improve animal versus human welfare. However, Wong notes that these implicit valuations are not themselves normatively plausible as the correct way to make trade-offs between humans and animals, and provide no guide to the correct trade-off between animal and human well-being from a non-anthropocentric point of view that sees animal well-being as fundamentally valuable in its own right.

So, the challenging question remains—how can we estimate the well-being capacity of species relative to humans? This question is not about the formal structure of the correct analysis, but is about its content; for example, how do we estimate the well-being level of an average bird versus the well-being level of an average human? This is the difficult problem that needs to be solved, and the theoretical formalism discussed earlier does nothing to help us solve it, although it does help in focusing our attention on the substantive question that needs to be answered.

Overcoming the Problem of Interspecies Comparisons

It is important to see that the problem of interspecies comparisons is not identical to and does not reduce to the fully general and familiar question of how to make interpersonal comparisons between human individuals, as we already have well-accepted methods for making interpersonal comparisons, for example, based on proxies for human well-being such as consumption,[23] human development indicators, and the like.[24] These empirical proxies for human well-being are generally assumed to ground good-enough estimates of individual human well-being levels for use in decision analyses via an assumed-to-be uniform relationship across individuals from the proxy for well-being (e.g., consumption) to the estimated level of well-being for each individual.

In other words, in economics, the challenge of making interpersonal comparisons is familiar. But it is also familiar how this problem is solved in practice, namely, by simply making comparisons based on a method that is believed to involve a good approximation: for example, by estimating a uniform concave function mapping consumption c to utility/well-being, such as

2. $$W^{TU} \approx \sum_{i \in humans} \frac{(c_i)^{1-\theta}}{1-\theta},$$

where θ parameterizes the diminishing marginal utility of consumption.

Typically economic practice is more crude and approximate than this, because economists often use population-level average consumption $\bar{c}$ as the proxy for the consumption of every individual, despite known inequality in individual consumption. For instance, the following equation simply multiplies the utility of per capita average consumption by the size of the human population P_h:

3. $$W^{TU} \approx P_h \frac{(\bar{c})^{1-\theta}}{1-\theta}$$

Mathematical formalism aside, the important point is merely that these formulas take consumption to be a proxy for well-being, and involve a concave transformation of consumption into well-being, which yields diminishing marginal utility of consumption where the rate of diminishment is parameterized by the θ term—which means that different views about the relationship between wealth and well-being can be investigated by varying the θ term. For example, if $\theta > 2$, then additional wealth above some sufficiency threshold generates little increase in well-being.

Our proposal for solving the problem of interspecies comparisons is analogous to the method used in Equations 2 and 3 of taking consumption as a proxy for human well-being: we propose to make interspecies comparisons based on a proxy that is imperfect but yet is as good as is possible in practice. To do this we first need a proxy, call it n, to use as the basis for estimating well-being potentials across species, analogous to the earlier use of consumption (c) as the basis for estimating well-being across humans. As an overly simplistic illustration of this idea (that might nevertheless be useful in practice in some contexts), n might be the number of neurons in the brain of members of a species. Data on number of neurons are readily available and may be a good proxy for well-being potential in some select contexts, such as an enormous global analysis involving billions of individuals where different species are crudely lumped together in small number of bins such as 'mammals' and 'insects;' this is true even if neurons are not the best proxy when fine-grained accuracy between individuals is more important.[25] For example, when greater accuracy is required for specific species or individuals, researchers can set n equal to a more complex metric based on expert analysis of empirical properties that are best correlated with different levels of cognitive capacity and hedonic enjoyment—for example, the number of neocortex-like neurons, cortisol levels, sociality, or other leading factors identified by the scientific community and philosophers as most closely correlated with the capacity to have complex thoughts and feelings, and whatever other empirical properties are found to be necessary for experiencing well-being.[26] At the limit, this sort of metric can reflect the true relationship between empirically measurable facts about individuals and their well-being capacity.

Abstracting for now from those details that are not essential to the core challenge of how to make interspecies comparisons, the first step of the proposal here is to parameterize an empirical proxy n with an exponential weight ψ into comparative well-being capacity for different species. The second step is to multiply this estimate of well-being capacity by a descriptive measure of the degree to which this potential is actually realized, f (i.e., the quality-of-life-adjustment term from Equation 1), to yield the desired well-being estimates:

4. $$W^{TU} \approx \sum_{is} n_{is}^{\psi} f_{is} \delta_{is}$$

(In ordinary language, the total amount of well-being is approximately equal to the sum over all individuals across species of that individual's empirical basis for well-being capacity, raised to the normative exponent [which determines the relationship between the empirical proxy and well-being capacity], multiplied by the flourishing level of that individual relative to its species capacity, multiplied by the duration of that individual's life.)

In practice, it is often more convenient to use species-level averages (where averages are denoted by a bar over the letter) as the proxy for well-being potential, which can then be multiplied by the species population P_s:

5. $$W^{TU} \approx \sum_{s} P_s \bar{n}_s^{\psi} \bar{f}_s \bar{\delta}_s$$

Equations 4 and 5 summarize the basis for our proposed method for making interspecies comparisons. They require an empirical proxy for n (e.g., number of neurons or a more complex empirically based metric), values for ψ grounded in normative and empirical considerations (on analogy with how values for θ in Equations 2 and 3 are grounded in normative and empirical considerations), and empirically determined values for f (based on empirical facts about how well members of the species are actually doing relative to their species potential).

With this method in hand, a decision analysis (e.g., between competing investment possibilities or between alternative public policies) can make use of a sensitivity test that investigates how optimal policy is sensitive to the normative parameter ψ that, in our earlier proposal, can be used to generate different estimates of the comparative well-being capacity of species. In other words, such a sensitivity test can use the preceding equations to capture the range of empirically grounded and principled estimates that represent normative uncertainty over how to estimate the well-being of animals of different species.

15031-3373e-2pass-r02.indd 97 11/15/2019 12:46:16 AM

Table 7.1 summarizes a standard sensitivity test of this type that illustrates how a sensitivity test that could be incorporated into global policy analyses (which usually only represent animals crudely based on a very small number of classifications). It relies on different principled ways of using the parameter ψ to estimate the potential well-being of a species as a function of the average number of neurons n in a member of that species.

Each estimate is expressed in terms of the well-being capacity of one human life-year ($\bar{n}_h^\psi$, for short), and thus, each estimate divides by the estimated well-being capacity of one human life. Estimate $1 = \frac{\bar{n}_s^\psi}{\bar{n}_h^\psi}$, with ψ set equal to 1 (a higher estimate of the capacity of animals), whereas estimate $2 = \frac{\bar{n}_s^\psi}{\bar{n}_h^\psi}$, with ψ set equal to 2 (a lower estimate of the capacity of animals). Other estimates are possible, for example, based on different empirical proxies n and other factors, such as different normative parameters and different views about a possible threshold cognitive capacity that arguably is necessary for well-being; see Budolfson and Spears (2019b) for more examples and discussion.

Each estimate can be used to put human life-years (which can be estimated via familiar proxies such as Equations 2 or 3) on the same scale as the life-years of animals of different species, and each estimate does so in a principled way that is empirically grounded. For example, assuming the number of neurons as a basis for well-being estimates, ifψ is set equal to 2 (a principled lower value for animals), then a human life-year is worth almost 120,000 mammal life-years and almost 120,000,000 fish life-years. If instead ψ is set equal to 1 (a principled higher value for animals), then a human life-year is worth about 344 mammal life-years and about 10,700 fish life-years. These alternative estimates appear to represent much of the range of empirically grounded and principled views over the well-being of animals of different species,[27] and can avoid unintuitive implications.[28] It may not even be desirable to attempt to choose between these estimates in policy analysis, if the goal is to take normative uncertainty into account and test the sensitivity of optimal decisions to this range of different reasonable (and empirically and theoretically principled) estimates.

In sum, the method developed in this section allows interspecies comparisons, via Equations 4 and 5, based on both empirically available estimates of species population dynamics and levels of flourishing for members of species[29] and also empirical proxies for well-being capacity n that can be calibrated with the ψ parameter to reflect normative uncertainty about the connection between those empirical proxies and well-being capacity. The key advance is that the term $\bar{n}_s^\psi$ provides a tool for making principled and empirically based estimates of the well-being capacity of different species of animals (i.e., $\bar{n}_s^\psi$ provides an estimate of the key uncertain term $\bar{\pi}_s$ in Equation 1). Thus, $\bar{n}_s^\psi$ provides a framework for articulating principled substantive answers to the question of how to make interspecies comparisons, and allows us to parameterize these comparisons to the range of normative uncertainty about their true value.[30]

Table 7.1 Two alternative estimates of the well-being potential of animal life-years of different species based on the number of neurons in an average member of the species, for illustrative purposes

Wildlife	*n*	*Utility Potential Estimate*	
	Number of neurons (n)	*(ψ = l)*	*(ψ = 2)*
Mammals	250	0.002907	0.000008450514
Birds	150	0.001744	0.000003042185
Amphibians etc.	15	0.000174	0.000000030422
Fish etc.	8	0.000093	0.000000008653
Insects etc.	0.1	0.000001	0.000000000001
Humans	86,000	1	1

Conclusion

This handbook chapter has explained the problem of interspecies comparisons. It has also explained recent research on developing methods to overcoming this problem, making it possible to include animal welfare in rigorous policy and investment analysis (e.g., in analyses of optimal public policies, analyses of optimal philanthropic investment, and so on). The development of these methods is important: methods of incorporating animal well-being will have an important impact on estimates of optimal prosocial investments of time and money by charities, businesses, or individuals, and similarly for estimates of optimal public policies for correcting market failures that ignore the costs of goods not reflected in their market price, for sustainable intensification of agriculture that aims to take animal welfare into account, for climate change, and for wilderness protection and other challenges related to natural resource usage.

Notes

1. Here and in what follows we use *animals* as shorthand for 'nonhuman animals.'
2. Budolfson 2015.
3. See Berkey forthcoming; Thomas forthcoming; Pacelle 2017.
4. See Wong 2016, GiveWell, Open Philanthropy Project, Animal Charity Evaluators.
5. Cowen 2006; Norwood and Lusk 2011 chapter 10; Jarvis and Donoso 2018.
6. Garnett et al. 2013; Norwood and Lusk 2011.
7. Hsiung and Sunstein 2007; Budolfson and Spears 2019a.
8. Hsiung and Sunstein 2007; Sunstein and Nussbaum 2004; Sunstein 2018 chapter 6.
9. Merely for ease of discussion, we frame all the discussion in this chapter in terms of consequentialism, noting here that deontological views can be represented as consequentialist views in the sort of analyses that are the focus of our discussion here (see Dreier 2011). In saying this, we do not take a stand on whether deontological views are adequately represented at the most fundamental level by 'consequentializing' them, we merely note that consequentialized versions are extensionally equivalent to the fundamental deontological views, and so deontological views can be adequately represented extensionally in the sort of analyses we are interested in here.
10. Technically, the interesting form of suboptimality here is Pareto-inferiority; namely, there is a way of making the outcome better for some that makes the outcome worse for no one.
11. Before concluding from this that therefore government regulations should be enacted to make the outcome better, it must be taken into account that sometimes new regulations would themselves be predictably inefficient, and as a result, the actual consequence of new regulations might in some cases predictably make things worse. See Budolfson 2017 and the references therein for discussion of this important complication for what normative public policy conclusions actually follow or don't follow from the widespread existence of market failure.
12. See Cowen 2006; Norwood and Lusk 2011 chapters 9 and 10 for extended argument for this.
13. See, for example, Norwood and Lusk 2011 chapters 9 and 10.
14. Jarvis and Donoso 2018; Otte and Chilonda 2000.
15. In contrast, when the anthropocentric value of animals is well reflected in market prices—such as, for example, the price of pollination services—market prices are the preferred method of valuation, at least to the extent that the good is a private good traded in a well-functioning marketplace.
16. See Hsiung and Sunstein 2007 and the references therein.
17. Ibid.
18. Schmidtz and Shahar 2018; Sandler 2018; McShane 2018; Palmer et al. 2014; Sarkar 2012; Jamieson 2008; O'Neill et al. 2008; Ng 1995; Singer 1975.
19. See references in previous footnote, and in addition Chan et al. 2016; Frank and Schlenker 2016; Dasgupta 2014; Alcamo 2003; Costanza et al. 1997; Kagan 2019.
20. See the relevant footnote earlier in this chapter on consequentializing moral theories. As an illustration, one can imagine a policy analysis that assumes deontological side constraints and then maximizes welfare subject to those constraints—the methods that follow are equally essential to such calculations as to unconstrained welfare maximization.
21. Compare the term $\overline{f}_s$ to McMahan 2001's concept of *fortune*, a connection Wong 2016 notes.
22. Estimating the flourishing term *f* can be seen as the focus of existing animal welfare science—see, for example, Fraser 2008; Appleby et al. 2011.

23. Consumption is often understood as income minus savings.
24. For an overview, see Adler and Fleurbaey 2016.
25. See Budolfson and Spears 2019a; Herculano-Houzel 2017; Olkowicz et al. 2016.
26. Herculano-Houzel 2017; Olkowicz et al. 2016; Barron and Klein 2016; Shriver 2014; Dawkins 2012; Appleby et al. 2011; Fraser 2008; see also Tye 2017; Persad 2020; Sebo 2020; Browning 2019; Fischer 2016.
27. Compare Alexander 2019.
28. See Budolfson and Spears 2019b.
29. See for example Fraser 2008; Appleby et al. 2011.
30. See Budolfson and Spears 2019b.

Bibliography

Adler, M., and Fleurbaey, M. (2016) *Oxford Handbook of Well-Being and Public Policy*, Oxford University Press.

Alcamo, J. (2003) *Ecosystems and Human Well-Being: A Framework for Assessment*, Washington, DC: Island Press.

Alexander, S. (2019) *Update to Partial Retraction of Animal Value and Neuron Number*, https://slatestarcodex.com/2019/05/01/update-to-partial-retraction-of-animal-value-and-neuron-number/.

Animal Charity Evaluators, www.animalcharityevaluators.org/.

Appleby, M., Joy Mench, I., Olsson, A., and Galindo, F. (eds.) (2011) *Animal Welfare*, 2nd ed., CABI.

Barnhill, A., Budolfson, M., and Doggett, T. (eds.) (2016) *Food, Ethics, and Society*, Oxford: Oxford University Press.

Barron, A., and Klein, C. (2016) "What insects can tell us about the origins of consciousness," *Proceedings of the National Academy of Sciences* 113(18): 4900–4908.

E[illegible]y, B. (forthcoming) "Prospects for an Animal-Friendly Business Ethics," in N. Thomas (ed.) *Animals and Business Ethics*, Springer.

Browning, H. (2019). *If I Could Talk to the Animals: Measuring Subjective Animal Welfare*, PhD Dissertation, Australian National University.

Budolfson, M. (2015) "Consumer ethics, harm footprints, and the empirical dimension of food choices," in A. Chignell, T. Cuneo, and M. Halteman (eds.) *Philosophy Comes to Dinner*, Routledge.

Budolfson, M. (2017) "Market failure, the tragedy of the commons, and default libertarianism in contemporary economics and policy," in D. Schmidtz and C. Pavel (eds.) *Oxford Handbook of Freedom*, Oxford University Press.

Budolfson, M., and Spears, D. (2018) *Why the Repugnant Conclusion is Inescapable*, Princeton CFI Working Paper.

Budolfson, M., and Spears, D. (2019a) *Optimal Climate Policy Including Animal Welfare*, Princeton CFI Working Paper.

~~Budolfson, M. and Spears, D. (2019b). *Public Policy, Consequentialism, the Environment, and Non-Human Animals*, in Portmore ed. Oxford Handbook of Consequentialism, Oxford UP~~.

Chan, K. et al. (2016) "Why protect nature? Rethinking values and the environment," *Proceedings of the National Academy of Sciences* 113(6): 1462–1465.

Costanza, R. et al. (1997) "The value of the world's ecosystem services and natural capital," *Nature* 387: 253–260.

Cowen, T. (2006) "Market failure for the treatment of animals," *Society* 43(2): 39–44.

Dasgupta, P. (2014) "Measuring the wealth of nations," *Annual Review of Resource Economics* 6: 17–31.

Dawkins, M. (19[illegible]) *Animal Suffering*, Springer.

Dawkins, M. (2012) *Why Animals Matter: Animal Consciousness, Animal Welfare, and Human Well-being*, Oxford: Oxford University Press.

Dreier, J. (2011) "In defense of consequentializing," *Oxford Studies in Normative Ethics*, Vol. 1. Oxford University Press.

Fischer, B. (2016). "What if Klein & Barron are right about insect sentience?", *Animal Sentience* ([illegible]).

Frank, E., and Schlenker, W. (2016) "Balancing economic and ecological goals," *Science* 353: 651–652.

Fraser, D. (2008) *Understanding Animal Welfare: The Science in its Cultural Context*, Wiley.

Garnett, T. et al. (2013) "Sustainable intensification in agriculture: Premises and policies," *Science* 341(6141): 33–34.

GiveWell, www.givewell.org/.

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

Hsiung, W., and Sunstein, C. (2007) "Climate change and animals," ~~155~~ *University of Pennsylvania Law Review* 1695.

Jamieson, D. (2008) *Ethics and the Environment*, Cambridge: Cambridge University Press.

Jarvis, L., and Donoso, P. (2018) "A selective review of the economic analysis of animal health management," *Journal of Agricultural Economics* 69(1): 201–225.

Kagan, S. (2019). *How to Count Animals, ~~m~~ore or ~~l~~ess*. Oxford ~~UP~~.

Klein, C., and Barron, A. (2016a) "Insect consciousness: Commitments, conflicts and consequences," *Animal Sentience* 9(21).

Klein, C., and Barron, A. (2016b) "Reply to Adamo, Key et al., and schilling and cruse: Crawling around the hard problem of consciousness," *Proceedings of the National Academy of Sciences* 113(27): E3814–E3815.

McMahan, J. (2001) *The Ethics of Killing*, Oxford University Press.

McShane, K. (2018) "Why animal welfare is not biodiversity, ecosystem services, or human welfare: Toward a more complete assessment of climate impacts," *Les Ateliers de l'Éthique/The Ethics Forum* 13(1): 43–64.

Mendl, M., and Paul, E. (2016) "Bee happy: Bumblebees show decision-making that reflects emotion-like states," *Science* 353: 1499–1500.

Muehlhauser, L. (2017) *2017 Report on Consciousness and Moral Patienthood*, Open Philanthropy Project.

Ng, Y-K. (1995) "Towards welfare biology: Evolutionary economics of animal consciousness and suffering," *Biology and Philosophy* 10: 255–285.

Ng, Y-K. (2016) "How welfare biology and commonsense may help to reduce animal suffering," *Animal Sentience*. http://animalstudiesrepository.org/cgi/viewcontent.cgi?article=1012&context=animsent

Norwood, F. B., and Lusk, J. L. (2011) *Compassion by the Pound: The Economics of Farm Animal Welfare*, Oxford: Oxford University Press.

Olkowicz, S. et al. (2016) "Birds have primate-like numbers of neurons in the forebrain," *Proceedings of the National Academy of Sciences* 113(26): 7255–7260.

O'Neill, J., Holland, A., and Light, A. (2008) *Environmental Values*, New York: Routledge.

Open Philanthropy Project, Farm Animal Focus, www.openphilanthropy.org/focus/us-policy/farm-animal-welfare.

Otte, M., and Chilonda, P. (2000) *Animal Health Economics: An Introduction*, Livestock Information, Rome, Italy: Sector Analysis and Policy Branch, Animal Production and Health Division (AGA), FAO.

Pacelle, W. (2017) *The Humane Economy: How Innovators and Enlightened Consumers Are Transforming the Lives of Animals*, William Morrow.

Palmer, C., McShane, K., and Sandler, R. (2014) "Environmental ethics," *Annual Review of Environment and Resources* 39: 419–442.

Persad, G. (2020) Cost-Effectiveness in Animal Health: An Ethical Analysis, in Bob Fischer ed. *Routledge Handbook of Animal Ethics*, Routledge.

Rushton, J. et al. (2018) "Initiation of global burden of animal diseases programme," *Lancet* 392(10147): 538–540.

Sandler, R. (2018) *Environmental Ethics: Theory in Practice*, Oxford: Oxford University Press.

Sarkar, S. (2012) *Environmental Philosophy: From Theory to Practice*, Wiley-Blackwell.

Schmidtz, D., and Shahar, D. (eds.) (2018) *Environmental Ethics*, 3rd ed., Oxford: Oxford University Press.

Sebo, J. (forthcoming) "Animals and climate change," in M. Budolfson, T. McPherson, and D. Plunkett (eds.) *Philosophy and Climate Change*, Oxford University Press.

~~Sebo, J. (2018). The Moral Problem of Other Minds, *The Harvard Review of Philosophy*.~~

Sebo, J. (2020). Effective Animal Advocacy, in Bob Fischer ed. *Routledge Handbook of Animal Ethics,* Routledge.

Singer, P. (1975) *Animal Liberation*, HarperCollins.

Singer, P. (2016) *Are insects conscious? Project Syndicate*, www.project-syndicate.org/commentary/are-insects-conscious-by-peter-singer-2016-05.

Shriver, A. (2014) "The asymmetrical contributions of pleasure and pain to subjective well-being," *Review of Philosophy and Psychology* 5: 135–153.

Sunstein, C., and Nussbaum, M. (eds.) (2004) *Animal Rights: Current Debates and New Directions*, Oxford: Oxford University Press.

Sunstein, C. (2018) *The Cost-Benefit Revolution*, MIT Press.

Tannsjo, T. (2016) "It's getting better all the time," in A. Barnhill, M. Budolfson, and T. Doggett (eds.) *Food, Animals and Society*, Oxford: Oxford University Press.

Thomas, N. (ed.) (forthcoming) *Animals and Business Ethics*, Springer.

Tooley, M. (1972) "Abortion and infanticide," *Philosophy & Public Affairs* 2: 37–65.

Tye, M. (2017) *Tense Bees and Shell-Shocked Crabs: Are Animals Conscious?* Oxford: Oxford University Press.

Wong, K. (2016) *Counting Animals: On Effective Altruism and the Prospect of Interspecies Commensurability*, B.A. thesis, Princeton University, Princeton, NJ.