Animal Cognition and Human Values[*]

Jonathan Birch

London School of Economics and Political Science
Email: j.birch2@lse.ac.uk
Webpage: personal.lse.ac.uk/birchj1

13 October 2017

Animal welfare scientists face an acute version of the problem of inductive risk, since they must choose whether or not to affirm attributions of mental states to animals in advisory contexts, knowing that their decisions hold significant consequences for animal welfare. In such contexts, the burden of proof should be sensitive to the moral consequences of error, but a framework for setting appropriate burdens of proof is lacking. Through reflection on two cases—the case of pain, and the case of cognitive enrichment—I arrive at a tentative general framework based on the principle of expected welfare maximization. I then discuss the limitations of this framework and the important questions it leaves open.

1. Inductive risk in the science of animal welfare

The interface of animal cognition research and animal welfare policy presents an important context in which to examine the role that ethical value-judgments should, or should not, play in the scientific process. The scientists who work at this interface—animal welfare scientists—have a dual role. On the one hand, they are practising scientists, seeking to advance our knowledge of animal welfare. On the other hand, they serve as expert advisors, shaping policy in ways that promote improved welfare for domesticated and captive animals. Good animal welfare policy should be sensitive to the needs and interests of animals; but, since animals cannot speak for themselves, the task of representing their needs and interests to policymakers usually falls to animal welfare experts.[1]

For a concrete example, consider the 2010 European Union (EU) directive on the protection of animals used for scientific purposes, which extended to cephalopod molluscs (such as

---

[*] This paper is based on my contribution to the PSA Women's Caucus Prize Symposium on "Animal Cognition and Animal Welfare", held at the 25th biennial meeting of the Philosophy of Science Association in Atlanta, GA, November 3-6, 2016. I thank my fellow symposium participants (Marta Halina, Kristin Andrews, Colin Allen and Lori Marino), I thank the PSA Women's Caucus for their support, and I thank the audience for their excellent questions and comments. I also thank Caroline Birch, Andrew Buskell, Marta Halina and an anonymous reviewer for their comments on the manuscript. This work was supported by a Philip Leverhulme Prize from the Leverhulme Trust.

[1] See, for example, the work of Marian Stamp Dawkins (1980) and Donald M. Broom (Broom 2014; Broom and Fraser 2015; Broom and Johnson 1993). On the importance of representing the needs and interests of animals in the policymaking process, see Kitcher (2015).

octopuses, squid and cuttlefish) the same regulatory framework that previously applied only to vertebrates. This directive was based on a 2005 report prepared by an advisory board of animal welfare scientists, drawing on the latest animal cognition research (AHAW 2005). The report had recommended extending protection to cephalopod molluscs and decapod crustaceans (such as crabs, lobsters and crayfish). However, following consultation with the EU member states, the recommendation regarding cephalopods was implemented, but the recommendation regarding decapods was not.

Because the animal welfare scientist is, and should be, both a scientist and a policy advisor, matters of ethical concern enter into the practice of animal welfare science in several ways. First, animal welfare scientists face an acute version of the problem of "inductive risk" (Rudner 1953; Douglas 2000, 2009), since they must choose whether to affirm or reject uncertain hypotheses about the mental capacities of animals, knowing that their decisions may hold significant consequences for animal welfare. Second, they face the challenge of transforming value-laden concepts, such as welfare, suffering and personhood, into objective and empirically measurable quantities. Third, they face questions of how to revise their own practices in the light of the ethical consequences of the discoveries they make.

This paper focusses primarily on the first issue—the problem of inductive risk. My concern is with the version of that problem that arises in the specific context of attributing mental states to animals. In general, one takes an inductive risk when one affirms or rejects a hypothesis that is uncertain, creating a risk of error. It is common to distinguish two types of error: a "Type I error" (or "false positive"), which involves the incorrect rejection of a true null hypothesis, and a "Type II error" (or "false negative"), which involves a failure to reject a false null hypothesis. These terms originated in the context of a particular statistical methodology—significance testing—but they have come to acquire a broader meaning, and are often applied in contexts where the null hypothesis is characterized qualitatively rather than statistically (the modern usage of the terms "false positive" and "false negative" is especially broad). I consider this broader, more informal usage harmless and will make use of it here.

In general, scientists tend to prioritize the avoidance of false positives, on the grounds that erroneously affirming a falsehood is worse than failing to affirm a truth. I suspect this stems from a general virtue of epistemic caution, or epistemic modesty; it is epistemically virtuous to avoid affirming falsehoods. In animal cognition research, the (qualitative) null hypothesis is usually the absence of the mental phenomenon of interest in the species of interest. False positives are thus errors of *over-attribution*, whereas false negatives are errors of *under-attribution*. Animal cognition researchers, in line with the epistemic caution of scientists in general, tend to prioritize the avoidance of errors of over-attribution. This is an idea famously captured in Lloyd Morgan's canon, a principle that has long been a source of debate and controversy among philosophers of biology (Morgan 1894; Allen-Hermanson 2005; Sober 2000, 2005; Fitzpatrick 2008; Andrews and Huss 2014).

Errors of over-attribution are often described as errors of anthropomorphism, and Morgan's canon is often interpreted as an imperative to prioritize avoiding such errors. More recently, others have described errors of under-attribution as errors of "anthropodenial" (De Waal 1999) or, more elegantly, "anthropectomy" (Andrews and Huss 2014). I avoid all of this terminology here; I will talk simply of over-attribution and under-attribution. Over-attribution may involve anthropomorphism but need not do so, and there may be forms of anthropomorphism other than over-attribution (see Section 6). Likewise, under-attribution may involve anthropectomy but need not do so, and anthropectomy too may come in other forms.

Perhaps there are contexts far removed from policy applications in which it is generally appropriate to set the burden of proof so as to prioritize avoiding errors of over-attribution. I take no stand on this. But when there are clear policy applications in view, I contend that it is not appropriate. For I contend that, when animal cognition research directly informs animal welfare regulations, a special context is created in which erroneously affirming a false mental state attribution may be a less serious error, all things considered, than failing to affirm a true attribution. The decision of whether to affirm a mental state attribution for the purpose of formulating animal welfare policy requires consideration of the moral consequences of error, and in this sense animal welfare science resembles the various policy-related areas of science discussed by Douglas (2009).

For this reason, and in contrast to animal cognition researchers working in less applied areas, few animal welfare scientists explicitly endorse Morgan's canon. The prevailing view is rather that an appropriate balance must be struck between tolerance of errors of over-attribution and tolerance of errors of under-attribution. However, a systematic framework for setting appropriate burdens of proof in animal welfare science is lacking. My aim in the rest of this paper is to take some initial steps towards the construction of such a framework by reflecting on two case studies: the case of pain, and the case of cognitive enrichment.

2. Risks of under-attribution: the case of pain

The full taxonomic range of the capacity to feel pain is unknown. Fiercely contested cases include fish (including bony, cartilaginous and jawless fish), arthropods (including crustaceans and insects) and molluscs (including cephalopods and gastropods). For each of these taxa, one finds some evidence in favour of pain, but one also finds sceptical critiques of the evidence. I will not survey the empirical literature on this issue here, rich and fascinating though it is (e.g. Bateson 1991; Allen 2004, 2013; Sneddon et al. 2014; Adamo 2016). I simply want to argue that this is a case in which the role of ethical considerations in setting an appropriate burden of proof seems relatively straightforward. For I take the following to be uncontroversial: to formulate animal welfare regulations on the assumption that animals of species *S do not* feel pain, when in fact animals of species *S do* feel pain, creates a risk of serious negative consequences for animal welfare.

For example, to exempt a cephalopod species from regulations regarding scientific experimentation, on the grounds that cephalopods do not feel pain, creates a risk of serious negative consequences for their welfare if it is in fact the case that they do feel pain. This is because it would allow these animals to be subjected, legally, to prolonged and intense pain, such as the pain caused by the removal of limbs without anaesthetic. I do not need to commit to any particular definition of welfare to make this claim, because I take it to be a constraint on any reasonable definition that it will respect the obvious relationship between prolonged, intense pain and negative welfare (the question of how to define welfare will be touched upon, briefly, in Section 6). By contrast, I contend that to formulate animal welfare regulations for *S* on the assumption that organisms belonging to *S do* feel pain, when in fact they *do not*, would *not* create a risk of serious negative consequences for animal welfare.

Given this asymmetry of risk between under-attribution and over-attribution, it is appropriate for animal welfare scientists to prioritize the avoidance of errors of under-attribution when attributing pain, at least when advising on animal welfare policy. One attractive way to do this is to adopt the following principle:

> (1) In the context of advising on animal welfare policy, an animal welfare scientist should affirm the hypothesis that pain is felt by organisms of species *S* whenever there is credible scientific evidence that pain is felt by organisms of *S*, even if that evidence is inconclusive and subject to continuing debate.

A number of animal welfare scientists have advocated versions of this view, and have noted its close relationship to the "precautionary principle" (Bradshaw 1998; P. Andrews 2011; Sneddon et al. 2014). For example, credible scientific evidence in this context might take the form of experiments showing the *self-delivery of analgesics*, whereby the animal learns to administer pain relief drugs such as opioids in an operant conditioning setup; *motivational trade-offs*, whereby the animal behaves as if weighing its preference to avoid pain against other preferences; or *conditioned place avoidance*, whereby the animal learns to avoid locations at which it previously encountered noxious stimuli (Sneddon et al. 2014). No one would argue that these behaviours conclusively indicate pain; but they make it credible that pain is experienced.

In many cases, however, the above principle is arguably not "precautionary" enough, because the number of species in contested taxa for which we have any relevant evidence at all is remarkably small. I have argued in previous work that a more practical approach is to take the *order* rather than the species as the appropriate level of analysis, and to affirm the hypothesis that pain is felt by organisms of a given order *O* whenever there is credible scientific evidence that pain is felt by organisms of *any* species within *O* (Birch 2017). For example, evidence of pain in a single species of decapod crustacean should be deemed sufficient for extending protection to the entire order of decapod crustaceans: we should not seek separate evidence for each of the approximately 15,000 species of that order.

Here I want to consider the extent to which this reasoning generalizes to other mental states. Setting aside the question of whether the species is the right level of analysis, one might hope to generalize Principle (1) by adopting a principle of the following form:

> (2) For any mental state *M*, under-attributing *M* creates far more serious risks of negative animal welfare outcomes than over-attributing *M*. So, in the context of advising on animal welfare policy, an animal welfare scientist should affirm the hypothesis that organisms of species *S* have *M* whenever there is credible scientific evidence that organisms of *S* have *M*, even if that evidence is inconclusive and subject to continuing debate.

In broad terms, Principle (2) is the inverse of Morgan's canon: *when in doubt, err on the side of over-attributing mental states*. I suggest, however, that although this general anti-canon may seem attractive at first sight, it oversimplifies a complex issue. In reality, the relative seriousness of over- and under-attribution depends on the species, mental state and animal welfare intervention in question. Consideration of a second case—the case of cognitive enrichment—will help us see why.

3. Risks of over-attribution: the case of cognitive enrichment

An important tool for promoting the welfare of captive or domesticated animals is *environmental enrichment*, which aims at creating opportunities for such animals to express behaviours they would express in the wild, stimulating brain activity and alleviating boredom. A subset of environmental enrichments, known in the animal welfare literature as *cognitive enrichments*, exploit animal cognition research so as to provide animals with challenges that allow them to exercise their cognitive abilities more fully, and that enable them to exert control over aspects of their environment (Meehan and Mench 2007).

What counts as a cognitive enrichment for a particular animal depends on the cognitive capacities of that animal, and there must be a close match between cognitive ability and environmental design. A challenge that is so difficult as to induce stress or anxiety is not an enrichment; a challenge so easy as to induce boredom or apathy is no enrichment either (Meehan and Mench 2007). Thus animal cognition research feeds directly into the design of appropriate cognitive enrichments.

This creates risks of over-attribution. For example, there is tentative evidence that cows, when presented with simple operant conditioning tasks, enjoy the experience of learning new skills. A study by Hagen and Broom (2004) involved measuring the heart rate of cows as they confronted a task that involved pushing a button to open a gate, leading to food. Learning how to solve the puzzle for the first time, or how to solve it more quickly, was predictive of increased heart rate, an effect that was not present in cattle who had already learned how to solve the puzzle, or in those who had not yet solved it. This, Hagen and Broom argue, provides some evidence that cows can recognize when they have learned something new, and

moreover—since heart rate is taken to be an indicator of affective arousal—some evidence that this experience alleviates boredom. Similar studies have been carried out on pigs and goats, with similar results (Puppe et al. 2007; Langbein et al. 2009; Zebunke et al. 2011). This indicates that operant conditioning tasks *may* constitute a valuable form of cognitive enrichment for farm animals (Manteuffel et al. 2009).

Yet, for many reasons, this evidence is clearly inconclusive. The cause of increased heart rate may have been something other than a recognition that the puzzle had been solved. Increased heart rate may not have been indicative of affective arousal at all. The affective arousal in question may have had a negative rather than a positive valence: that is, the animals might have felt frustrated or anxious rather than excited and pleased (there were no other signs that this was the case, but the evidence does not rule it out). The individuals or breeds in the sample may not be representative of other individuals or other breeds. The sample may simply have been too small. Jumping too quickly to the conclusion that operant conditioning tasks generate positive welfare, and recommending this form of welfare intervention prematurely, creates a risk of causing unnecessary stress to individuals who cannot solve the puzzles, or who gain no benefit from solving them. There might also be opportunity costs, if ineffectual environmental enrichments were to be recommended at the expense of simpler, more effective ones.

Hence it is appropriate here to proceed with caution, and to require more than the mere existence of credible scientific evidence that cows enjoy puzzle solving before advising the implementation, in farms or zoos, of enrichments premised on their possession of that ability. Hagen and Broom (2004, 212) implicitly acknowledge this, writing that "because of the novelty of the approach and the small number of animals, this study should be seen as a first step towards further investigation of the topic", despite quoting a high level of statistical significance (a p-value of 0.009) on the key result.

In general, it seems appropriate to require more than the mere existence of credible scientific evidence that an organism possesses a cognitive ability before recommending an enrichment premised on its possession of that ability. The context of designing enrichments to alleviate boredom is very different from that of devising regulations to minimize animal pain, and a different burden of proof is appropriate. We should therefore reject Principle (2).

4. Varying the burden of proof: a starting point

Can a general framework accommodate the considerations raised by the foregoing cases? I will now sketch, tentatively, what such a framework might look like, before proceeding to discuss some limitations and unresolved issues. An attractive framework for setting burdens of proof, I suggest, is an expected welfare maximization framework based on the following overarching principle: the burden of proof for affirming a mental state attribution, in the context of advising on animal welfare policy, should be set so as to maximize the expected total welfare of the non-human animals affected by the policy.

Why include only *non-human* animals? Why do the consequences of error for human well-being not merit consideration when setting the burden of proof? This exclusion is justified on the grounds that the role of the animal welfare expert in advisory contexts is to act as a representative for the needs and interests of non-human animals. In a sound policymaking process, the needs and interests of humans are already well represented by policymakers and representatives of relevant sectors. If animal welfare scientists were also to factor in effects on humans, these effects would be double-counted. For example, in the consultative process that led to the 2010 EU directive on the protection of animals used for scientific purposes, the UK Bioscience Sector (2009) pushed hard for various concessions, including the exclusion of decapods. If the AHAW (2005) panel had already taken due account of their concerns when judging whether to affirm that decapods feel pain, these concerns would have been double-counted. At this stage in the process, it was appropriate to ignore them.

Given this overarching principle, we can describe the animal welfare scientist's decision problem as follows. Suppose an animal welfare scientist X must choose, in a policy context P, whether or not to affirm the hypothesis that species S has mental state M. This might, for example, be the hypothesis that a cephalopod or decapod species feels pain, in the context of formulating a set of regulations for laboratory research. There are four possible outcomes (Table 1): a correct attribution (A), a correct non-attribution (D), an error of over-attribution (C) or an error of under-attribution (B).[2]

	X affirms in P that S has M	X does not affirm in P that S has M.
S has M	A	B
S lacks M	C	D

(Table 1)

Let W represent the total welfare of the non-human animals affected by the policy to be formulated in P. The introduction of such a variable presupposes that we can measure welfare quantitatively, and that there is a way of aggregating the welfare of many individual animals, which may mean aggregating over individuals of many different species (e.g. all cephalopods). These presuppositions merit substantive debate in their own right; some might object that comparing welfare across species is to compare apples with oranges. But I am presupposing the possibility of an overall welfare measure here in order to think about how we should set the burden of proof for mental state attributions, on the assumption that there exists such a measure (I revisit some of the unresolved issues here in the final section).

---

[2] This approach is inspired by Jeffrey's (1956) decision-theoretic reconstruction of Rudner (1953), as well as by Steele's (2012) more recent reconstruction.

By the principle of expected welfare maximization, X should accept in P the hypothesis that S has M if and only if:

E (W | K & X accepts in P that S has M) > E (W | K & X does not accept in P that S has M)

Here K denotes the scientist's relevant background knowledge, and the E ("expectation") operator takes a weighted sum of the possible welfare outcomes, with each outcome weighted by its conditional probability given the scientist's background knowledge and her decision as to whether or not affirm the hypothesis.

On the assumption that X's affirmation or non-affirmation of the hypothesis makes no difference to the conditional probability (relative to K) that the hypothesis is correct, this can be rearranged to yield the following, potentially more useful criterion (where Pr denotes probability):

Pr (S has M | K) / Pr (S lacks M | K) > [W(C) − W(D)] / [W(B) − W(A)]

The above criterion shows how the burden of proof should vary, in a context-sensitive way, with the relative seriousness of a false positive and a false negative. For every S, M and P, there is some critical probability that warrants affirming in P that S has M, but the critical probability depends on the details of the case. For example, when M is pain, S is an animal widely used in laboratory research, and P is the formulation of regulations for such research, the critical probability will be low, because an error of under-attribution is likely to be far more serious than an error of over-attribution. When M is the capacity to recognize that one has learned a new skill, S is a farm animal, and P is the recommendation of a cognitive enrichment that would yield no benefit in the absence of M, the critical probability will be substantially higher.

5. Jeffrey's objection
In the 1950s literature on inductive risk, Jeffrey's (1956) well-known objection to Rudner (1953) was that most scientists, in practice, are not in a position to assign utilities to outcomes such as A, B, C and D, because the wider social consequences of affirming or rejecting hypotheses are usually unforeseeable. Moreover, in cases where they are foreseeable, it seems as though lawmakers and voters—not unaccountable scientists—should make the relevant value-judgements. We can see now why animal welfare science constitutes a special context in which these objections are not compelling.

First, we can give greater substance to the notion of utility by defining it, in this context, as the total welfare of the affected non-human animals. As noted above, this brackets the difficult issue of how welfare is to be measured, and how cross-species comparisons are to be made. However, since the field of animal welfare science is premised on the assumption that

welfare can indeed be measured, this is not an unreasonable presupposition for my purposes (though see Section 6).

Second, by noting that affirmation is not affirmation *simpliciter* but affirmation *in a specific advisory context,* we can single out a particular causal path from a mental state attribution to welfare consequences that makes it feasible to evaluate the comparative severity of over-attribution and under-attribution, relative to that context. In some cases, such as the case of pain, the welfare consequences of error are readily foreseeable. In other cases, such as the case of cognitive enrichment, the welfare consequences of error are complex and open to debate, sometimes requiring further investigation in their own right, but they are far from wholly unforeseeable, and they can and should be investigated empirically.

Third, we can see that it is animal welfare scientists, and not politicians or voters, who are typically best placed to evaluate the seriousness of the different types of error. In other contexts, the individuals affected by a policy decision, or their elected representatives, may be better placed to make the relevant ethical value-judgements (cf. Kitcher 2002, 2011). But the context of animal welfare science is one in which the affected individuals cannot speak for themselves (Kitcher 2015). It falls to the animal welfare scientist, qua policy advisor, to give due weight to negative welfare consequences of over- or under-attribution.

6. Limitations, open questions and future directions
I suspect animal welfare scientists tacitly recognize many of these points; the role for philosophers of science is to clarify, systematize and defend the role that value-judgements already play in this field. But I hope to have shown that this is not a trivial task. Reflection on the case of pain shows very clearly that ethical concerns are relevant to the burden of proof, but reflection on other cases shows the relationship to be more subtle than it may first appear. The above sketch of an expected welfare maximization framework provides a tentative way forward, but it also draws our attention to two important open questions.

First, there is the problem of defining, quantifying and aggregating animal welfare. There are obvious connections here to the problem of defining, quantifying and aggregating human well-being in the philosophy of psychology (see, e.g., Alexandrova 2012). In both cases, we confront similar challenges. How do we respect the need to make animal welfare objective and empirically tractable while also doing justice to its normative value, and to its subjective, first-person character? Animal welfare scientists often invoke the concept of *coping*: good welfare is said to consist in successful coping with the environment, where coping involves "control of mental and bodily stability" (Broom and Johnson 1993). Defenders of such accounts maintain that coping is, on the one hand, empirically measurable, while being, on the other hand, a normatively valuable property that captures welfare's subjective and internal aspects. But there is room for further debate here, and these qualitative definitions still leave open the question of how welfare should be quantified in particular cases, and how cross-species comparisons should be drawn.

Second, there is the problem of understanding how the cognitive capacities of an animal causally influence its welfare. The case of pain is an easy case in this respect: it is obvious that pain bears negatively on welfare. The case of cognitive enrichment is a harder case. It is intuitive to suppose that engagement in a stimulating task is conducive to good welfare, whereas the stress of failing a task or the boredom of completing an unchallenging task are not. Yet there is a danger here of a form of *evaluative anthropomorphism*—a projection of our own values on to animals—that is distinct from the cognitive anthropomorphism often involved in over-attributing mental capacities. Might it be, for example, that stress can sometimes generate welfare benefits in captive animals, due to its role in facilitating learning? This suggestion (discussed by Meehan and Mench 2007) may seem counterintuitive, but it should not be dismissed out of hand. More generally, as we turn to more complex mental capacities studied by comparative psychologists—such as theory of mind, long-term memory, causal reasoning and reasoning about the future—the likely welfare implications of attributing (or misattributing) these capacities become increasingly difficult to discern, in the absence of a richer theoretical understanding of the nature and causes of good psychological welfare in animals.

In sum, an expected welfare maximization framework tells us how a theoretical understanding of the relationship between an animal's cognitive capacities and its welfare should inform judgements of the burden of proof for mental state attributions, in the context of formulating animal welfare policy. But I freely acknowledge that this leaves us with much of the hard work—the work of better understanding the causal pathways linking cognition to welfare—still to do. The imperative to maximize expected welfare provides an abstract framework for setting appropriate burdens of proof for the attribution of mental states, *provided we understand how those mental states, if present, affect welfare*. But with some exceptions, such as the case of pain, that understanding is currently lacking.

This short paper should not, therefore, be regarded as solving the problems it raises. Its main contribution is to show why a deeper theoretical understanding of the relationship between animal cognition and animal welfare is urgently needed, if we are to set appropriate burdens of proof for mental state attributions to animals in policymaking contexts.

References

Adamo, Shelley Anne. 2016. "Do Insects Feel Pain? A Question at the Intersection of Animal Behaviour, Philosophy and Robotics." *Animal Behaviour* 118:75-79.

Alexandrova, Anna. 2012. "Well-Being as an Object of Science." *Philosophy of Science* 79:678-89.

Allen, Colin. 2004. "Animal Pain." *Noûs* 38:617-43.

-----. 2013. "Fish Cognition and Consciousness". *Journal of Agricultural and Environmental Ethics* 26:25-39.

Allen-Hermanson, Sean (2005). Morgan's Canon Revisited. *Philosophy of Science* 72:608-31.

Animal Health and Animal Welfare Panel of the European Food Standards Agency (AHAW). 2005. "Aspects of the Biology and Welfare of Animals Used for Experimental and Other Scientific Purposes." *The EFSA Journal* 292:1-46.

Andrews, Kristin and Brian Huss. 2014. "Anthropomorphism, Anthropectomy, and the Null Hypothesis." *Biology and Philosophy* 29:711-29.

Andrews, Paul L. R. 2011. "Laboratory Invertebrates: Only Spineless, or Spineless and Painless?" *ILAR Journal* 52:121-5.

Bateson, Patrick. 1991. "Assessment of Pain in Animals." *Animal Behaviour* 42:827-39.

Birch, Jonathan. 2017. "Animal Sentience and the Precautionary Principle," *Animal Sentience* 2(16). http://animalstudiesrepository.org/animsent/vol2/iss16/1/

Bioscience Sector. 2009. *A Bioscience Sector response to the House of Lords European Union Committee Inquiry: Revision of Directive 86/609 on the protection of animals used for scientific purposes*.

Bradshaw, R. Harry. 1998. "Consciousness in Non-Human Animals: Adopting the Precautionary Principle." *Journal of Consciousness Studies* 5:108-14.

Broom, Donald M. 2014. *Sentience and Animal Welfare*. Wallingford: CABI.

----- and A. F. Fraser. 2015. *Domestic Animal Behaviour and Welfare* (5th edition). Wallingford: CABI.

----- and K. G. Johnson. 1993. *Stress and Animal Welfare*. Dordrecht: Kluwer.

Dawkins, Marian Stamp. 1980. *Animal Suffering: The Science of Animal Welfare*. London: Chapman and Hall.

De Waal, Frans B. M. 1999. "Anthropomorphism and Anthropodenial: Consistency in Our Thinking about Humans and Other Animals." *Philosophical Topics* 27:255-80.

Douglas, Heather E. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67:559-79.

-----. 2009. *Science, Policy, and the Value Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press.

Fitzpatrick, Simon. 2008. "Doing Away with Morgan's Canon." *Mind and Language* 23:224-246.

Hagen, Kristin and Donald M. Broom. 2004. "Emotional Responses to Learning in Cattle." *Applied Animal Behaviour Science* 85:203-13.

Jeffrey, Richard C. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science* 23:237-46.

Kitcher, Philip. 2002. *Science, Truth and Democracy*. New York: Oxford University Press.

-----. 2011. *Science in a Democratic Society*. Amherst, NY: Prometheus Books.

-----. 2015. "Experimental Animals". *Philosophy and Public Affairs* 43:287-311.

Langbein, Jan, Katrin Siebert and Gerd Nürnberg. 2009. "On the Use of an Automated Learning Device by Group-Housed Dwarf Goats: Do Goats Seek Cognitive Challenges?" *Applied Animal Behaviour Science* 120:150-58.

Manteuffel, Gerhard, Jan Langbein and Birger Puppe. 2007. "Increasing Farm Animal Welfare by Positively Motivated Instrumental Behaviour." *Applied Animal Behaviour Science* 118:191-98.

Meehan, Cheryl L. and Joy A. Mench. 2007. "The Challenge of Challenge: Can Problem Solving Opportunities Enhance Animal Welfare?" *Applied Animal Behaviour Science* 102:246-61.

Morgan, C. Lloyd. 1894. *An Introduction to Comparative Psychology*. London: The Walter Scott Publishing Company.

Rudner, Richard. 1953. "The Scientist qua Scientist Makes Value Judgments". *Philosophy of Science* 20:1-6.

Sneddon, Lynne U., Robert W. Elwood, Shelley A. Adamo and Matthew C. Leach. 2014. "Defining and Assessing Animal Pain." *Animal Behaviour* 97:201-12.

Sober, Elliott. 2000. "Evolution and the Problem of Other Minds." *Journal of Philosophy* 97:365-387.

-----. 2005. "Comparative Psychology Meets Evolutionary Biology." In Lorraine Daston and Gregg Mitman (eds.), *Thinking with Animals*. New York: Columbia University Press, pp. 85-99.

Steele, Katie. 2012. "The Scientist qua Policy Advisor Makes Value Judgments." *Philosophy of Science* 79:893-904.

Zebunke, Manuela, Jan Langbein, Gerhard Manteuffel, Birger Puppe. 2011. "Autonomic Reactions Indicating Positive Affect during Acoustic Reward Learning in Pigs." *Animal Behaviour* 81:481-89.