

The Six Principles, Philosophy, and Applying Human Ethics to Animals

Julian Savulescu

The Six Principles

Beauchamp and DeGrazia's (BD's) Six Principles are arguably the most constructive step forward in the ethics of animal experimentation in the last 50 years. BD aim to develop and significantly augment Russell and Burch's Principles of Humane Experimental Technique, which involved the so-called 3 Rs: replacing sentient animals with other models where possible; reducing the number of animal subjects to what is needed for statistical adequacy; and refining techniques to reduce animal suffering. Their reason for moving beyond the 3 Rs is that "this framework does not feature general moral principles that display the core values at work in animal research ethics. In addition, the 3-Rs framework does not adequately address the costs and benefits of animal research to human beings or include a comprehensive program of animal-subjects protection" (p. ...).

I want to further these aims of the Six Principles. As BD correctly observe, "the core values of animal research ethics are social benefit and animal welfare." The critical issue is how to weigh human interests against animal interests. This is, in practice, where a great fudge can loom large: either the human benefit is overestimated or the harm to animals underestimated. I want to address how these benefits and harms can be more precisely evaluated and how they should be balanced.

In the first section, I aim to explicate further Principle 3 (The Principle of Sufficient Value to Justify Harm) and Principle 5 (The Principle of Basic Needs), though I will often assume rather than cite or appeal to these principles. In the second section, I will

explicate Principle 6: The Principle of Upper Limits to Harm. Basic needs will be especially prominent in my discussions. However, my points are often sufficiently general that they pertain to all of the Six Principles: there are value judgements in interpreting and specifying all these principles. We need new ways to think about fleshing them out. The most innocuous of the principles are 1 (No Alternative Method) and 4 (The Principle of No Unnecessary Harm), but even in these cases, we will need to specify what counts as an "alternative" and "unnecessary." There might be some alternative method, but it might be difficult to implement. There might be some way of avoiding pain but it might be wasteful of resources and have significant opportunity costs. What is really meant is "no reasonable alternative" and we then need a way of deciding what is reasonable, in which case we are back to weighing human and animal interests.

This problem, of course, does not escape BD. It is clearly expressed when they write, "The Principle of Basic Needs requires investigators and caretakers to meet animal subjects' basic needs *unless failure to meet one or more basic needs is an unavoidable consequence of, and morally justified by,* scientific procedures in a well-designed protocol that satisfies the three principles of social benefit" (p. •••, italics mine). One example of successful specification or "precisification" is BD's list of basic needs. This list makes clear what must be addressed by researchers and ethics committees. It supplies a plausible catalogue of the basic needs that includes the following:

- “• Nutritious food and clean water
- Safe shelter
- Adequate stimulation, exercise, and opportunities for species-typical functioning
- Sufficient rest to maintain physical and (where applicable) mental health
- Veterinary care

- For social species, access to compatible conspecifics or social group members
- Freedom from significant experiential harms such as pain, distress, and suffering
- Freedom from disease, injury, and disability
- Freedom of movement with adequate space

It is controversial whether the following is a basic need:

- Freedom from premature death" (p. ●●●).

My procedure hereafter will be as follows: after commenting on the evolution of research ethics from a “big picture” perspective, I will turn to address the issue of the badness of premature death in animals. I will then try to show both how these animals' interests in avoiding premature death can be weighed against human interests and how we can set an Upper Limit to Harm (following Principle 6). These matters are critical to questions of the practicability of the BD principles. Since BD do not specifically consider novel animals who have different needs—those who have been created by new reproductive technologies or by gene editing to have different needs—I will also address the ethics of creating such animals, including animals who do not feel pain or suffer.

In what follows, I will use two well developed theoretical ideas from human ethics and apply them to animal ethics:

1. The veil of ignorance. This procedure was developed by John Rawls to identify what would constitute a just and fair state of affairs for humans. It involves imagining oneself behind a veil or curtain, not knowing who in society one would be, whether one would be the best off, the average person, or the worst off. Rawls famously claimed that from behind the veil, one would choose the social and legal arrangements which made the worst off as well off as possible (maximin).¹
2. The distinction between two different kinds of reasons: "person"-affecting and "impersonal" reasons. Person-affecting reasons involve harm or benefit to

particular human beings, while impersonal reasons involve making the world more generally being better or worse, without there being harm or benefit to any particular individual.²

I will support the idea that ethics is universalizable, including ethical judgments that start from human ethics but also apply to animals.

Big Picture: The Evolution of Research Ethics

As BD briefly observe, the *Belmont Report* was a major advance in the ethical oversight of human research. Beauchamp was a driver of that report. It was and remains a great achievement that was published by the National Commission for the Protection of Human Subjects at a particularly timely point (1978) in the early history of bioethics. In this report (as elsewhere in bioethics), Human Ethics Committees, or Institutional Review Boards (IRBs) as they are called in the U.S., were given responsibility for ethical oversight of research involving humans and were situated at the center of the process of approving research protocols. Today IRBs flourish and tightly control research throughout the world. However, it is time to reinvent ethics review to address the ways technology, ethics, and science have evolved. As I have argued elsewhere, ethics committees now often obstruct valid research with lethal consequences,³ and often fail to properly execute their ethical obligations.⁴

We now need a second revolution in both human research ethics and animal research ethics. I have argued that the key element of human ethics is ensuring that the risks involved in research are "reasonable" rather than minimal. The five criteria I have outlined for reasonable risk in humans are remarkably similar to those BD outline for animals. In determining whether the risks of participation in research are reasonable, the following factors are relevant:⁵

1. Is there a known risk to participants prior to commencing the study, and what is its magnitude, based on evidence available at the time?
2. Should any non-human or epidemiological research, systematic overview, or computer modelling be performed prior to the study to better estimate the risk to participants or obviate the need for the use of human participants?
3. Could the risk have been reduced in any other way? Is it as small as possible?
4. Are the potential benefits (in terms of knowledge, improvement of welfare of trial participants or other people) of this study worth the risks?
5. Could this research generate knowledge that is likely to significantly harm either participants or others outside the research, now or in the future?

The same kinds of consideration can and should be applied to animal research ethics by substituting "animal" for "human": what is a reasonable risk to the animal subject or a reasonable harm and what is proper oversight of levels of risk?

In ethics, people often think in polar terms such as safe vs unsafe, proven vs unproven, and certain vs uncertain, but much of ethics is grey. Things are rarely entirely necessary or safe. There are costs and benefits as well as probabilities, and these have to be weighed across different populations such as current patients vs future patients⁶ and animals vs humans. We need new ways to think more robustly and deeply about these inescapably ethical and philosophical questions of fairness and justice, and BD do a great

job of delineating the arena. These are not mere scientific questions; they are fundamental moral questions.

BD point to the central role of animal research ethics committees. These are important, but they must evolve too and increase the professionalism and sophistication of their deliberations. I have argued elsewhere that their judgments should be objective, not merely subjective.⁷ The current major problem is that committee members have wide discretion to decide cases on their intuitions, including those about whether or not there is "sufficient scientific justification" or "sufficient protection of animal welfare." We should do better than that. We need principled ways of thinking about the Six Principles and instantiating them in practice. This commentary is an attempt to think about this problem.

The Reincarnation Test

Imagine that Indian religions—Buddhism, Hinduism, and Jainism—are right.

Reincarnation exists. After death, your spirit enters the ether. In this world of reincarnation, you could come back in the next life either as a different human being or as an animal. ("Animal" is used as shorthand for "non-human animal" in this comment.)

What system of animal ethics would you choose? The Reincarnation Test is a version of the veil of ignorance. Addressing this question through a thought-experiment about reincarnation can help achieve an impartial view of human uses of animals in research. I would balance the benefits of the use of animals in research now against the life I might lead as a rat, monkey, or other laboratory animal. Animal research might lead to medical breakthroughs that prolong my life or improve its quality now or in the future as a human. But it also might not, and in the reincarnation model I might come back as a monkey subjected to painful procedures, kept in a cage, with a much shortened life.

From this position behind the veil of ignorance,⁸ I would choose the various protections afforded by the Six Principles that BD have developed. I would want there to be a reasonable chance the research would lead to human benefit, and I would want animals to have a reasonably decent life. I want, however, to suggest a benchmark for what a reasonably decent life is for the animal, one that I believe is entirely consistent with the BD account. I want to do so using the reincarnation thought experiment and the philosophical model of a veil of ignorance for at least three reasons:

1. While some philosophers argue that animal suffering counts the same as human suffering,⁹ others argue that animal suffering counts for less than human suffering.¹⁰ Even if the latter is true, there are vastly more animals than humans. There are 7 billion humans. While there are fewer than 800 million vertebrate lab animals, there are more than 24,000 billion vertebrate farm animals,¹¹ a figure that does not include pets. The chances of coming back as an animal directly under human control are vastly greater than benefiting significantly in our lives from animal research or coming back as a human. If one considers all of the future, and if current trends continue, the sheer number of future animals under human control will be vast.

2. There is no good reason to apply different standards to research than to other areas of life. The ethics we derive for animals should apply not only to research, but to farming and the rearing of pets—a point about consistency in our thinking about animal ethics generally. The same basic standards of protection should apply to all animals whose lives are directly and intentionally controlled by humans.

3. Aristotle's Principle of Equality appropriately requires that we treat like cases alike unless there is a morally relevant difference.

I will now assume that these three reasons and the veil of ignorance idea are acceptable and will move on to examine whether the principles that should govern

imposing harm or risk of harm on humans for public interest or social reasons can be extended to animals.

The Veil of Ignorance, Lethal Experimentation, and Uncontrollable Epidemic

To take more seriously how we should treat animals, we should start with how we should treat humans. In the bioethics community we have thought far more about human ethics than about animal ethics, so we should try to learn how to generalize from human ethics to bring it to bear on animal research ethics. While all guidelines for medical research proscribe significant harm to humans, there are good reasons to question the ethical applicability of such proscriptions in all cases. Before considering research ethics I find it helpful to consider the conditions under which intentionally killing human beings in an extreme emergency might be morally justified.

Imagine an uncontrollable epidemic afflicts humanity. It is highly contagious and eventually every single human will be affected. It causes people to fall unconscious. Five out of six people never recover and die within days. One in six people mounts an effective immune response. They recover over several days and lead normal lives. Doctors can test people on the second day, while still unconscious, and determine whether they have mounted an effective antibody response or whether they are destined to die. There is only one treatment. Doctors can extract all the blood from the one in six people who do mount an effective antibody response on day 2, while they are still unconscious, and extract the antibodies. There will be enough antibodies to save five of those who don't mount responses, though the extraction procedure will kill the donor. The five will go on to lead normal lives, and the antibody protection will cover them for life.

If you were a person in Uncontrollable Epidemic, which policy would you vote for? The first policy, Inaction, is one in which nothing is done. One in six of the world's population survives. The second policy is Extraction, which kills one but saves five others. There is no way to predict who will be an antibody producer. You don't know if you will be one of the six who can mount an immune reaction or one of the five in six who don't manage to mount an immune response and would die without the antibody serum.

Put simply, you don't know whether you will be one who would survive or one who would die without treatment. All you know for certain is that you will catch the disease and fall unconscious. You may recover or you may die while unconscious. Inaction gives you a 1 in 6 chance of being a survivor. Extraction gives you a 5 in 6 chance.

This case is easy for many consequentialists. Extraction saves five times as many lives and should be adopted because it has the best overall consequences. But which would you choose behind the Rawlsian veil of ignorance, not knowing whether you would be immunocompetent or immunodeficient? I would choose Extraction. I would definitely become unconscious, like others, and then there would be a 5 in 6 chance of waking up to a normal life. This policy could also be endorsed on Kantian contractualist grounds. Not only would rational self-interest behind a veil of ignorance endorse it, but it could be willed as a universal law. Consequentialism and contractualism here converge.¹² I believe other moral theories would also endorse Extraction.

While in Uncontrollable Epidemic we can consent in principle to the possibility of being killed, animals typically cannot consent to research performed upon them. However, if we imagine the lives of the noncompetent animals and humans affected by a decision, the veil of ignorance serves to provide a kind of rational impartial consent. For

example, we could imagine what it would be like to be unconscious (nothing) and form a view about the interests of such an individual receiving life sustaining medical treatment.

Extending the Argument to Animals: A Fair Go for Animals

The reasoning above can be extended to both lethal experimentation and nonlethal but harmful experimentation on animals. Return to reincarnation. What kind of animal life would we choose, if we might be either a human or a lab animal? First, basic needs would have to be met, thereby satisfying the BD Principle of Basic Needs. What we might want is something like a fair innings or "fair go"¹³ for the animal. This would be a life which not merely avoided negatives - freedom from pain, hunger, etc. There would need to be positive features to life:

- socialization
- play
- instinctive behavior
- procreation and offspring rearing
- happiness, contentment

BD cover important headings of this sort in their list of Basic Needs, but we need more detail on what such a life would be like. For example, if I were to return as a chicken, I would minimally want:

1. A suitable social environment (no more than few hundred conspecifics)
2. An ability to consume and forage for a variety of food (not just a standard mash diet)
3. An ability to express the full variety of natural behaviors: dust bathe in suitable substrate (soil, fresh litter or sand; not in manure/litter mixtures); sun-bathe (UV lamps or sun); perch; lay eggs in a secluded nest; forage, etc.

4. Appropriate housing to provide ambient temperatures and shelter from predators.
5. Getting killed painlessly before becoming seriously sick or injured.¹⁴

Let's assume this chicken is in all morally relevant respects similar to research animals with comparable moral status. How long should such a research animal live? Here we can employ the Reincarnation Test. How long would we want to live as a research animal if we did not know whether we were a human who might benefit from research or an animal whose life would be shortened? What I would want as a chicken is a decent life. A decent minimum for a chicken (or "fair innings" as it has been termed in human ethics - sometimes taken to be "three score and ten") might be to reach sexual maturity. By engaging in sexual behavior and procreating, such an animal would have achieved a major important life "milestone" for the species.

The degree of self-consciousness of animals also matters. The more self-conscious the animal, the more a longer life presumably matters because the animal has expectations, perhaps even plans, for the future. The more complex and rich one's mental life, hopes, dreams, expectations, and memories, the more tragic it would be for that life to be cut short.¹⁵

If an animal lacks self-consciousness, less is lost in premature death. While it might be a decent life for a chicken to live long enough to reach, say, sexual maturity, if chickens lack significant self-consciousness, then it might not be seriously wrong to end the chicken's life before that point for some especially important scientific research, though it may not be justifiable for the production of food. Thus we might believe that the criteria listed above ought to be met while a chicken is alive, but it could be ethical to kill it painlessly before that point for some pressing need. However, for great apes it would be wrong to shorten their lives before the decent minimum because an ape, unlike

a chicken, is significantly self-conscious, so the loss imposed on the ape is greater than that imposed on the chicken.

For purposes of animal research ethics, it seems reasonable to assume that if we bring a sentient being into existence to be a research subject, we should only do so if we can give it a minimally decent life while it is alive, and self-conscious animals should have a fair innings. For example, if one chooses to have a pet, that pet should have a minimally decent life as judged by the Principle of Basic Needs. It need not have the very best chance of the best life. A sufficiently good life is enough.

Mindless and Radically Genetically Modified Animals: Person-Affecting vs Impersonal Reasons

In this section, I will give a philosophical explication of BD's final principle: the Principle of Upper Limits to Harm. This somewhat different subject about research involving animals is about radical genetic modification, such as creating animals without minds. In his chapter "The Moral Status of the Cloning of Humans,"¹⁶ Michael Tooley raised the possibility of creating cloned anencephalic human infants as a source of organs for transplantation. This approach is extendable to research. Gene editing of gametes or embryos could be used to create mindless anencephalic human infants who could be used for experimentation, perhaps sparing animals who can experience pain from experimentation.

Similar proposals have been made to create animals that do not feel pain, as a way of reducing suffering.¹⁷ In part, breeding has already achieved these kinds of modifications: some breeds of chickens have been bred to not care about their eggs.¹⁸ In order to ethically evaluate such proposals for creating modified animals for research, we can employ a distinction drawn by Derek Parfit in human ethics between modifications

that are identity-altering (that is, they affect which or whether an individual comes into existence) and modifications that are identity-preserving (that is, they affect an already-existing individual for better or worse). To return to the human example, gene editing an embryo to be deaf would be like deafening a future child, say by cutting its auditory nerves. This would be person-affecting (and therefore identity-preserving) and so more wrong than selecting a human embryo that was created from eggs or sperm with an already existing and unavoidable mutation for deafness. This might be impersonally wrong, but less wrong than causing person-affecting harm.¹⁹ In general, many philosophers hold that person-affecting wrongs are worse than impersonal wrongs.²⁰

Similarly, genetically altering an already-existing animal not to care about its young or not feeling pain is identity-preserving, harmful, and *prima facie* wrong. (I am here assuming an animal incapable of experiencing pain will overall lead a worse life, for example, by injuring itself or by facilitating more crowded confinement.) It is wrong to gene edit animals to suffer from disease—unless we follow the qualifications in BD’s moral framework that allow certain harms to research animals when necessary for and morally justified by scientific purposes. Thereby the reason for the wrongness of gene editing to make an animal’s life worse might be outweighed by reasons of scientific and human benefit.

We might think that using animals having characteristics that lower their quality of life tied to their identity (that is, determine their identity) is not wrong. For example, it seems that using an animal that could not otherwise exist except for some genetic disposition that lowers its quality of life but makes it more useful to humans is not wrong. A mule is bred by mixing donkey sperm and a horse egg to work hard, so imposing hard work on it might not be thought wrong. However, this assessment would be a mistake. Even if it is more wrong to deafen a human embryo than it is to select a naturally deaf

embryo, it would be wrong not to gene edit or otherwise correct deafness in a naturally deaf embryo. Indeed, I believe it would be as wrong as it is to deafen an embryo. So too the amount of work a mule should be subjected to is governed by the Six Principles, not by the fact it was bred to work hard. When we have the power to improve human lives, there are strong reasons of beneficence to improve them and of nonmaleficence not to harm them. And so too for animals. If we could gene edit out the tendency of hens bred not to care about their own eggs, we should do so, if it would make their lives go better (other things being equal).

The impersonal vs animal-affecting distinction does not allow us to create or use animals with lower prospects of well-being when we can use technology like gene editing to improve their lives. Creating animals to not feel pain or even using naturally occurring breeds (where this feature could be corrected) would not get us off the hook. Our obligations to animals stand.

Suppose a genetic modification would alter identity. For example, genetically modifying an embryo so that the animal would never be conscious would be identity altering (on some accounts of personal identity—given the assumption, which I find plausible, that the capacity for consciousness creates a new individual with a mind). This modification would be akin to preventing the animal with consciousness from coming into existence. It would be like abortion or embryo destruction. This would not be wrong assuming abortion or embryo destruction is permissible. Creating such a mindless animal would be like creating in vitro meat for human consumption.²¹

To summarize thus far, while it would be wrong to modify animals who can or would experience pain to not experience pain, it would not be wrong to create unconscious animals from scratch. And it may not be necessary to go all the way to unconsciousness. For example, gene editing a pig embryo (before the animal comes into

existence with interests) so that the resulting animal was so severely retarded that it had few interests would be morally akin to destroying the pig embryo and would be identity-altering. In such cases, one is not harming the animal (e.g., the pig) by lowering its capacity to experience pain. One is stopping the pig from coming into existence and replacing it with a different, severely retarded individual animal.

Such an account can be applied to non-self-conscious animals once they exist, for example chickens. To painlessly kill a chicken is permissible because it lacks a higher-order mental property such as self-consciousness. If an animal has mental states but minimal psychological continuity and connectedness—minimal self-awareness over time—then biological modification of the chicken to remove its consciousness is not wrong; it is akin to killing it. And this extends to other radical modifications that fracture the chain of identity. It is like creating a new animal. So, paradoxically while it is wrong to lower a chicken's capacity to experience pain (because it is harmful in the "chicken"-affecting sense), it may be less wrong to make that chicken unconscious or severely retarded. However, there may be impersonal reasons not to create such animals (for example, it would be better for the world if there were happy chickens rather than mindless chickens, even if it would not be better for any individual chicken).

To summarize this section, lowering the capacities for a good life in animals which do already or will exist (either through genetic selection or editing) is *prima facie* wrong and violates the Principle of Upper Limits to Harm by causing significant harm. However, certain modifications would not cross the threshold of the Principle of Upper Limits to Harm. Radically modifying the early stages of life of self-conscious animals prior to that animal achieving self-consciousness is not wrong on grounds of nonmaleficence if it is identity-altering, though it may be wrong on other impersonal grounds. Similarly, gene editing or other biological modification of characteristics of

non-self conscious animals is not wrong on animal-affecting grounds of harm if it is identity-altering (for example, by removing consciousness or creating severe cognitive impairment), though it may be wrong on other impersonal grounds.

Minimal Beneficence and Duty of Easy Rescue

According to Minimal Beneficence or a duty of easy rescue, where risks of performing an action are small for the rescuers, and the benefit (or harm averted) to others is large, we should perform that action. Thus when the benefit to animals of some act (such as supporting humane farming) is great, we should act in that way if the cost to humans is sufficiently small. The reverse is also true: perhaps we should live a little less long so that our research animals have much better lives, in accordance with fair innings arguments. Such considerations also imply that animal research might be most justified as preliminary studies for pediatric research rather than research into chronic conditions that affect people after their own fair innings.

The Principle of Minimal Beneficence is consistent with several of BD's Six Principles:

The Principle of Sufficient Value to Justify Harm—this principle would not be met when the benefit to humans is very small and the harm to animals large.

The Principle of Basic Needs—this principle will generally entail that harm to the animal is great if basic needs are not met.

The Principle of Upper Limits to Harm—this principle identifies the point at which the harm to the animal is great.

Few would deny Minimal Beneficence or a duty of easy rescue. This moral norm clearly converges with at least three of the Six Principles.

Conclusion

To sum up and conclude, in the context of scientific research we may have good reasons to bring animals into existence as research subjects especially when they are to be given high quality lives as judged by the BD Principle of Basic Needs. However, these reasons may be outweighed by animal-affecting considerations, and in practice, we may have only weak reasons, if any, to bring animals into existence for their own sake. If we do bring animals into existence, strong animal-affecting reasons require us to enable these animals to have a high probability of a sufficiently good life, although these reasons must be weighed in the context of scientific research against human-centered reasons. The facts may be such that some decrement in the *length* or the *quality* of the animal's life is justified for research on the basis of the BD principles of social benefit (assuming the BD principles of animal welfare are also satisfied) just as a decrement in a human being's length or quality of life could be justified in a context of research on human subjects by similar reasoning.²²

Our attitudes and practices towards animals are often wildly parochial and inconsistent. It is instructive to compare the enormous lengths people go to providing for their pets (even bequeathing them their entire fortune²³) to the treatment of lab animals all the way to factory farming, where little real consideration of animal welfare occurs. In a recent uproar in the U.S., a science teacher faced six months in jail for feeding a dying puppy to a pet tortoise in front of pupils.²⁴ Officials killed the tortoise. In Japan, dogs and cats enjoy special legal protections. Koreans have been denounced for farming and eating dogs, while we happily eat pigs. Some species of fish are protected, like the shark (which kills many fish), but not others. Our treatment of animals seems largely driven by whether they are seen as friendly, attractive, or special or useful to us. These drivers are

understandable, but they do not provide moral justifications of our preferences in animal research ethics.

In my arguments I have tried to show that BD are right to look to principles as basic to justification. BD have taken animal ethics to the next level and have given reasons why it is now time to further explicate and specify their Six Principles into operationalizable criteria for animal research—and for human-animal relations more generally. I have tried to show how philosophical approaches to dilemmas in human ethics—the veil of ignorance and the distinction between person-affecting and impersonal reasons—can help us to more precisely operationalize the Six Principles. Animal ethics is as complex and difficult as human ethics, though we don't yet realize it.

Around forty years ago Tom Beauchamp and James Childress revolutionized medical ethics by introducing the Four Principles of Biomedical Ethics²⁵—respect for autonomy, beneficence, nonmaleficence, and justice. These were intended to be "mid-level" principles between a variety of ethical theories and approaches and actual medical practice. In 25 years of practice as an academic practical ethicist, I have found these invaluable. However, they leave much open for rich ethical discourse: what is autonomy, how much beneficence is morally required (for example, is it a duty of easy rescue, or more), and what constitutes justice? In the same way, the Six Principles in the Beauchamp and DeGrazia *Principles* volume are best interpreted as the animal research equivalent: mid-level principles that leave open for many decades to come a rich body of insights that can and should inform our interpretative and elaborative discourse.

Acknowledgements

Thanks to David DeGrazia, Tom Beauchamp and Patrick Birkel for very helpful comments.

NOTES

¹ John Rawls, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, 1971)

² Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984).

³ Julian Savulescu, "Bioethics: Why Philosophy is Essential for Progress," *Journal of Medical Ethics* 41 (2015): 28-33.

⁴ Julian Savulescu, Iain Chalmers, and Jennifer Blunt, "Are Research Ethics Committees Behaving Unethically? Some Suggestions for Improving Performance and Accountability," *British Medical Journal* 313 (1996): 1390-1393; Julian Savulescu and Merle Spriggs, "The Hexamethonium Asthma Study and the Death of a Normal Volunteer in Research," *Journal of Medical Ethics* 28 (2002): 3-4; Julian Savulescu, "Harm, Ethics Committees and the Gene Therapy Death," *Journal of Medical Ethics* 27 (2001): 148-150.

⁵ Julian Savulescu, "Safety of Participants of Non-Therapeutic Research Must be Ensured," *British Medical Journal* 316 (1998): 891-2; Julian Savulescu and Tony Hope, "Ethics of Research," in J. Skorupski (ed.), *The Routledge Companion to Ethics*, (Abingdon: Routledge, 2010), pp. 781-795.

⁶ Julian Savulescu and Tony Hope, *op.cit.*, note 5.

⁷ *Op. cit.*, note 3.

⁸ John Rawls, *A Theory of Justice*.

⁹ Peter Singer, "Speciesism and Moral Status," *Metaphilosophy* 40 (2009): 567-581.

¹⁰ Shelly Kagan, *How to Count Animals, More or Less*. Uehiro Series in Practical Ethics (Oxford: Oxford University Press, forthcoming).

¹¹ Brian Tomasik, "How Many Wild Animals Are There?" *Essays on Reducing Suffering* (2009, as last edited 2018), available at <http://reducing-suffering.org/how-many-wild-animals-are-there/>, retrieved 5 October 2018.

¹² Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), volume 1.

-
- ¹³ John Leslie Mackie, “Rights, Utility and Universalization,” in R.G. Frey (ed.), *Utility and Right* (Minneapolis: University of Minnesota Press, 1984), pp. 86-105; Julian Savulescu. “How Do We Choose Which Life to Save? Equality of Access or a Fair Go?” *Current Paediatrics* 12 (2002): 487-492.
- ¹⁴ Patrick Birkel, Personal Communication.
- ¹⁵ Peter Singer, “Life's Uncertain Voyage,” in Philip Pettit, Richard Sylvan, and Jean Norman (eds.), *Metaphysics and Morality: Essays in Honour of J.J.C. Smart* (Oxford: Blackwell, 1987).
- ¹⁶ Michael Tooley “The Moral Status of the Cloning of Humans” in J.M. Humber and R. F. Almeder (eds.), *Human Cloning. Biomedical Ethics Reviews* (Totowa NJ: Humana Press, 1998).
- ¹⁷ Marcus Schultz-Bergin, “The Dignity of Diminished Animals: Species Norms and Engineering to Improve Welfare,” *Ethical Theory and Moral Practice* 20 (2017): 843-856.
- ¹⁸ Frederick B. Hutt, *Genetics of the Fowl* (New York: McGraw-Hill Book Company, Inc, 1949). As cited in M. N. Romanov, “Genetics of Broodiness in Poultry—a Review,” *Asian-Australasian Journal of Animal Science* 14 (2001): 1647-1654.
- ¹⁹ Julian Savulescu, “The Nature of the Moral Obligation to Select the Best Children,” in Akira Akabayashi (ed.), *Future of Bioethics: International Dialogues* (Oxford: Oxford University Press, 2014), pp. 170-182.
- ²⁰ J. L. Coleman and C. W. Morris, eds. *Rational Commitment and Social Justice: Essays for Gregory Kavka* (Cambridge, UK: Cambridge University Press, 1998).
- ²¹ Owen Schaefer and Julian Savulescu, “The Ethics of Producing In Vitro Meat,” *Journal of Applied Philosophy* 31 (2014): 188–202.

²² Julian Savulescu, “The Structure of Ethics Review: Expert Ethics Committees and the Challenge of Voluntary Research Euthanasia,” *Journal of Medical Ethics* 44 (2018): 491-493. <http://dx.doi.org/10.1136/medethics-2015-103183>.

²³ James Yeates and Julian Savulescu, “Companion Animal Ethics: A Special Area of Moral Theory and Practice?” *Ethical Theory and Moral Practice* 20 (2017): 357-359.

²⁴ Peter Stubley, “Teacher Who Fed Live Puppy to Turtle in Front of Pupils Charged with Animal Cruelty,” *Independent*, available at <https://www.independent.co.uk/news/world/americas/puppy-fed-snapping-turtle-school-teacher-robert-crosland-preston-junior-high-idaho-a8381146.html>, retrieved 5 October 2018.

²⁵ Tom L. Beauchamp and James F. Childress, *Principles of Biomedical Ethics*, 8th edition (New York: Oxford University Press, 2019). The first edition was published in 1979.