

Integrating Human and Nonhuman Research Ethics

Jeff Sebo

Abstract In this chapter I argue for developing a unified moral framework for assessing human and nonhuman subjects research. At present, our standards for human subjects research involve treating humans with respect, compassion, and justice, whereas our ethical standards for nonhuman subjects research merely involve (half-heartedly) aspiring to replace, reduce, and confine our use of nonhuman animals. This creates an unacceptable double standard and leads to pseudo-problems, for example regarding how to treat human-nonhuman chimeras. I discuss general features that a more integrated moral framework might have, assess the pros and cons of this kind of this framework, and suggest that the pros decisively outweigh the cons.

1. Introduction

At present, we accept radically different moral frameworks for assessing human and nonhuman subjects research. On one hand, we accept a very high standard for morally permissible human subjects research, according to which we are morally required to treat humans with respect, compassion, and justice. On the other hand, we accept a very low standard for morally permissible nonhuman subjects research, according to which we are merely required to replace, reduce, and refine our use of nonhuman animals to the extent that doing so is compatible with achieving our scientific goals. The result is that we categorically forbid a wide range of harmful, lethal, nontherapeutic, nonconsensual research on human subjects while generally permitting such research on nonhuman subjects.

We treat human and nonhuman subjects differently not only because we accept different moral frameworks for human and nonhuman subjects research, but also because our assessment of this research takes place in a cultural and institutional context that heavily favors humans over nonhumans. When all decision-makers are human, we can expect human bias to affect our application of moral principles. And when educational opportunities, occupational opportunities, and research teams, facilities, and technologies center on a particular kind of research – in this

case, nonhuman subjects research – we can expect these factors to affect our application of moral principles too. As a result, not only do we accept lower standards for our treatment of nonhuman subjects, but we also take less care in our application of these standards.

These different approaches to human and nonhuman subjects research constitute an unacceptable double standard, and they also create pseudo-problems in research ethics. In particular, these different approaches lead us to treat humans and nonhumans differently in ways that cannot plausibly be justified on the basis of species difference alone. As a result, they also lead us to ask ethical questions that we might otherwise be able to avoid. For example, many people wonder how we should treat human-nonhuman chimeras. Should we apply “human” standards to them and treat them very well, should we apply “nonhuman” standards to them and treat them very poorly, or should we strike a balance between these extremes? Yet this question would be nonsensical with a more principled, integrated approach to research ethics.

In this chapter, I discuss our current, highly fragmented moral frameworks for assessing human and nonhuman subjects research and make the case for a more integrated approach. I suggest that a more integrated approach would aspire to treat all sentient beings with respect, compassion, and justice, but would also allow us to treat different animals differently on the grounds that, for instance, the nature and strength of our interests and needs vary across species. This more integrated approach has benefits and risks; for instance, it makes our approach to research ethics more consistent and principled, but it also risks making it simplistic, reductive, and potentially alter the course of scientific and medical progress. But I suggest that the benefits decisively outweigh the risks, particularly if we work to mitigate the risks.

Before I begin, I should make a couple of caveats about my approach to this chapter. First, I will assume in this chapter that all sentient animals – that is, all animals who are capable of consciously experiencing positive and negative states like pleasure and pain – have moral status – that is, morally matter for their own sake, and so we morally ought to consider their interests and needs when deciding how to treat them.¹ I will also assume that, even if all sentient animals have moral status in this sense, we can still be justified in treating different animals differently if they have different interests, needs, or other morally relevant features.² Of course,

¹ For a classic consequentialist argument for this idea, see Singer 1975, and for a classic non-consequentialist argument for this idea, see Regan 1983.

² For an argument that we should treat different animals differently in light of their capacities, see Kagan 2019. For an argument that we should treat them differently in light of our relationships with them, see Palmer 2010.

not everyone will agree with these assumptions. But ethicists have spent decades defending them, and my aim here is to build on that work rather than replicate it.

Second, I will not assume a particular moral theory – such as utilitarianism, rights theory, virtue theory, or care theory – in this chapter, nor will I attempt to flesh out the details of a more integrated approach to human and nonhuman subjects research ethics. Instead, I will discuss the general features that I expect a more integrated approach to have, focusing on features that can serve as the basis for “overlapping consensus” among different moral theories.³ And I will assess the general benefits and risks of a more integrated approach in these terms. I will also discuss further questions that we need to answer as we flesh out the details, for instance concerning how to assess the nature and strength of nonhuman interests and rights. But I will not attempt to answer these further questions here.

2. The Human Paradigm

In general, the human subjects research paradigm assumes a rights-based moral framework. We assume that treating humans well involves more than simply increasing human happiness and decreasing human suffering in the aggregate. It also means treating humans with respect, compassion, and justice along the way, by respecting human autonomy and balancing the benefits and burdens that we impose on individuals and groups. As a result, this rights-based moral framework tends to prohibit research that would impose excessive burdens on vulnerable individuals or populations – particularly when they are not capable of providing informed consent and we are not capable of providing them with compensatory benefits – even when this research has the potential to produce valuable knowledge.⁴

Consider each of these points in turn. First, human subjects research tends to place a high premium on *respect* for research subjects. In particular, we aspire to respect the autonomy of research subjects as much as possible. When humans are capable of providing informed consent, we require that researchers secure informed consent. Otherwise we proceed by seeking informed consent from relevant third parties, as well as by seeking assent or dissent – that is, an expression of approval or disapproval with respect to particular interactions – from the research subjects.

³ For discussion of the idea of overlapping consensus, see Rawls 1987 and Fleischacker 2011.

⁴ For general discussion of the ethics of human subjects research, see Resnik 2018.

And in the case of particularly harmful or risky research, we simply decide not to proceed at all. We also set a high bar for consent, for example by treating humans as incapable of consent in cases where they face strong social or economic pressure to say “yes.”

Second, and relatedly, human subjects research tends to place a high premium on *compassion* for human research subjects. We generally limit how many harms and risks we can impose on human research subjects, even when they provide informed consent, and even when this harmful or risky approach is necessary for the science. We also generally allow research subjects to stop participating in research at any point, without retribution, and we allow flexibility in the timing or location of office visits, medicine administration, and other activities so that research participation can fit into a full life. While these accommodations might make research less efficient, they are necessary both as a matter of respect and compassion, since they allow for sustained consent from, and care for, research subjects.

Third, and also relatedly, human subjects research tends to place a high premium on *justice* for human research subjects. As noted above, we attempt to avoid conducting excessively harmful research on vulnerable humans, particularly when they lack the ability to provide informed consent and when they lack access to meaningful alternatives (Grady 2005). And when we do conduct harmful research on vulnerable humans, we attempt to compensate not only the research subjects but also their communities, for instance by making sure that resulting benefits are accessible to community members (Bracken-Roche et al 2017). Otherwise we risk a situation where the burdens of research flow disproportionately to the worst-off among us and the benefits of research flow disproportionately to the best-off among us.

In light of these considerations, we generally set a high bar for morally permissible human subjects research. In particular, we generally hold that this research is morally permissible only when it respects human autonomy, limits harm to individuals and communities, and provides compensatory benefits to individuals and communities. As a result, we prohibit many studies that might produce valuable knowledge on the grounds that they violate one or more of these duties. In short, when it would be impossible to (a) secure informed consent from potential research subjects or guardians such as parents, (b) limit the harms that we impose on individuals or communities, or (c) compensate individuals or communities for the harms that we impose on them, we generally decide not to proceed.

Indeed, one might even argue that the bar for morally permissible human subjects research is sometimes *too* high. For example, during the COVID-19 pandemic, many bioethicists called for human challenge trials that would expose consenting adults to a controlled dose of COVID-19 in a controlled setting rather than wait for them to catch any amount out in the world.⁵ But even though a wide range of humans volunteered to participate, and even though the expected benefits of this approach clearly outweighed the expected harms, our leaders have generally decided not to pursue human challenge trials because they felt that the risks were too high.⁶ Many other cases have similar features, raising the question whether we sometimes forbid human subjects research that we should have permitted instead.

3. The Nonhuman Paradigm

In contrast, the nonhuman subjects research paradigm assumes a welfare-based moral framework. We assume that nonhuman animals have welfare but not rights, and we also assume that human interests generally trump nonhuman interests. As a result, we do not aspire to treat nonhumans with respect, compassion, and justice in the same ways that we do with humans, particularly vulnerable humans. Instead, we aspire to follow the “three Rs.” That is, we aspire to replace, reduce, and refine our use of nonhumans where possible, while allowing ourselves to perform harmful, lethal, nonconsensual, nontherapeutic research on nonhumans where “necessary.” And we define ‘necessity’ simply in terms of what means are necessary to achieve a particular scientific or medical aim.⁷

Consider each of these points in turn. First, we do not aspire to treat nonhumans with respect, compassion, or justice in the same ways that we do with humans, particularly vulnerable humans. First, we do not seek informed consent from guardians and assent or dissent from research subjects nearly as much as we do with humans. Second, we do not aim to limit the harm that we cause nonhumans nearly as much as we do with humans, nor do we aim to compensate them for harm caused nearly as much as we do with humans. Third, we do not aim to distribute the benefits and burdens of research equitably across human and nonhuman populations. And in

⁵ For an open letter signed by many bioethicists, including me, see here: <https://www.1daysooner.org/us-open-letter>

⁶ The UK permitted small COVID-19 challenge trials to proceed in October 2020, but even this approval occurred nearly a year after vaccines were ready for testing (Callaway 2020).

⁷ For general discussion of the three Rs, see Russell et al 1959.

general, insofar as we consider human and nonhuman welfare in harm-benefit analyses, we tend to prioritize human welfare over nonhuman welfare.

Instead, we aspire to follow the three Rs in nonhuman subjects research. According to this method, when we evaluate proposed nonhuman subjects research, we start by asking: Can we achieve the same goals without using animals at all? If so, we should. If not, then we ask: Can we achieve the same goals while using fewer animals? If so, we should. Either way, we then ask: Can we achieve the same goals while harming each animal less? If so, then we should. If not, then we generally permit the harm. In theory, this method allows us to harm animals insofar as, and only insofar as, we need to do so in order to achieve our scientific or medical aims. In practice, this method allows us to conduct harmful and lethal research on an estimated 100+ million nonhumans per year (Taylor and Alverez 2020).

As this description makes clear, the nonhuman subjects research paradigm thus defines ‘necessity’ in terms of what means are necessary to achieve a particular scientific or medical end. As we have seen, in the human case we generally allow ethics to trump science, prohibiting harmful, lethal, nonconsensual, and nontherapeutic studies *whether or not* we see them as necessary for achieving scientific or medical ends. In contrast, in the nonhuman case we generally allow science to trump ethics, permitting such studies when we see them as necessary for achieving particular scientific or medical ends. Put differently, in the human case we remember that instrumental rationality requires *either* taking the means *or* giving up the end, but in the nonhuman case we tend to forget that we have the second option.⁸

In light of these considerations, we generally set a low bar for morally permissible nonhuman subjects research. We proceed on the assumption that a wide range of harmful, lethal, nonconsensual, nontherapeutic, animal studies are morally permissible on the grounds that they have the potential to produce “knowledge worth having,” and that no alternative methods currently available to us have the same potential.⁹ The result is that we currently breed, raise, harm, and kill millions of nonhuman animals per year in order to produce epistemic and social benefits for humans. And while we do provide (varying levels of) care to these animals, we still treat them in ways that we would never treat humans – particularly humans who are incapable of providing informed consent – in modern science and medicine.

⁸ For discussion of the idea of ‘necessity’ in nonhuman subjects research, see Ferrari 2019.

⁹ For discussion of the idea of knowledge worth having, see Egel et al 2020 and Sebo and DeGrazia 2020.

While we can debate whether or not the bar for morally permissible human subjects research is sometimes too high, there is no debating that the bar for morally permissible nonhuman subjects research is, in the vast majority of cases, far too low. We should treat humans and nonhumans alike with respect, compassion, and justice, taking into account both the similarities and differences across species. And if this is right, then we should discount nonhuman welfare much less than we do, harm nonhuman animals much less than we do, and benefit nonhuman animals much more than we do. In short, the only possible justification for the status quo is the view that nonhuman interests carry either no weight at all or vanishingly little weight, and this view is simply not morally acceptable.

4. The Cultural and Institutional Context

Unfortunately, the ethical gap between human and nonhuman subjects research oversight is even greater in practice than in theory, since our application of the principles of human subjects research ethics is much more rigorous than our application of the principles of nonhuman subjects research ethics, due to the cultural and institutional context of each kind of research. In particular, when everyone involved in the decision-making process is human, it can be easy to let that affect our decisions. And when our research infrastructure is built to support particular kinds of research, it can be easy to let that affect our decisions as well. The result is that we achieve respect, compassion, and justice for humans much more than we achieve replacement, reduction, and refinement for nonhumans.

For instance, consider how we use harm-benefit analysis in nonhuman subjects research. We think: On one hand, this study will harm many animals. On the other hand, this study has the potential to contribute to scientific progress, and if it does, then it has the potential to benefit *very* many humans. In other words, we reason that the expected benefits outweigh the expected harms, since even though the probability of benefit is very low, the level of benefit is high enough to compensate for that. And given the nature of science, it can be hard to predict which studies will contribute to scientific progress, and which contributions to scientific progress will, in turn, produce social benefits. As a result, this reasoning can seem persuasive with respect to a wide range of harmful, lethal studies (Sebo and DeGrazia 2020).

But notice how this application of harm-benefit analysis stacks the deck in favor of harming animals. We consider long-term benefits for humans via scientific progress, but not long-term risks for humans via false positives or negatives in toxicity or efficacy, long-term risks for nonhumans via normalization of exploitation and extermination, or opportunity costs for humans and nonhumans via neglect of animal-free alternatives. We also credit scientific progress to animal research without considering the counterfactuals, that is, without considering whether we might have produced the same benefits via animal-free alternatives, without producing the same costs. And we merely aim to do more good than harm rather than aiming to do as much good and, perhaps more importantly, as little harm as possible (Bass 2012).

Our application of the Three Rs is similar. When evaluating harmful, lethal nonhuman subjects research, we tend to decide that replacing the use of animals is impossible when we are unaware of, or unprepared for, animal-free alternatives. We also tend to set limits on how much we can reduce and refine our use of animals, since, for instance, we might need to use a particular number of animals for our findings to be valid, and we might need to forego many methods of reducing harm to animals or increasing support for animals because we think that these methods will undermine the science or because we see them as too expensive. For instance, we kill instead of retire the vast majority of lab animals partly to collect further data and partly because there are simply too many for us to affordably retire.

That we assess nonhuman subjects research in these ways is predictable given the cultural context of this work. Everyone involved in the decision-making process, from the researchers proposing the research, to the committees evaluating the research, to the policy-makers supporting this activity, to the community members electing the policy-makers, is human. We all have human interests and perspectives, and we are all vulnerable to self-interest, speciesism, status quo bias (that is, a bias in favor of the status quo), scope insensitivity (that is, insensitivity to the significance of numbers), and other human biases. We also make these decisions in a cultural context that assumes human exceptionalism and the moral permissibility of nonhuman exploitation and extermination for human purposes.

That we assess nonhuman subjects research in these ways is also predictable given the institutional context of this work. In some jurisdictions, people see animal research as part of the

best, if not the only, possible route to approval for new foods or drugs.¹⁰ Additionally, many research teams and facilities are built for animal research. These factors make people more likely to believe that animal research is ethically, legally, or, at least, practically necessary. The result is a kind of institutional path dependence, where everyone assumes that animal research is necessary, and senior scholars create educational and occupational opportunities for junior scholars accordingly. The “need” for animal research then becomes a self-fulfilling prophecy, made apparently true by our failure to invest in animal-free alternatives.

5. Double Standards and Pseudo-problems

As a result of these different moral frameworks, as well as these different cultural and institutional contexts, we currently have an unacceptable double standard in human and nonhuman subjects research ethics. We treat humans much better than we treat nonhumans, to a degree that cannot plausibly be justified by species differences alone. This double standard reveals inconsistency and creates pseudo-problems. For instance, many people wonder how we should treat human-nonhuman chimeras: Should we treat them well, like we treat humans, or should we treat them badly, like we treat nonhumans? But many of these questions arise only because our approach to research ethics is so unprincipled and inconsistent. They would disappear entirely with a more principled and consistent approach.

To be clear, my claim here is not that we should treat humans and nonhumans the same way. As Peter Singer famously argued, equal consideration is compatible with differential treatment (Singer 1975, Chapter 1). For instance, to the degree that humans and nonhumans have different interests and needs, we might have different moral duties to them accordingly. To the degree that some animals have stronger interests and needs than others, we might have stronger moral duties to the former animals accordingly. Depending on which moral theory we accept, we might also think that we have different, or stronger, duties to some animals than to others in light of our relationships other features of our context; for instance, we might think that we owe more to animals we have harmed than to animals we have not (Palmer 2010).

¹⁰ For example, see this page from the United States Food and Drug Administration website: <https://www.fda.gov/emergency-preparedness-and-response/mcm-regulatory-science/animal-rule-information>

So when I say that we currently have an unacceptable double standard in human and nonhuman research ethics, I am not merely saying that we treat humans and nonhumans differently. I am saying that these differences cannot plausibly be justified by differences in interests, needs, histories, relationships, or other morally relevant features. For instance, even if we accept that we should generally prioritize humans on the grounds that humans generally have stronger interests and needs and we generally have stronger histories and relationships with them, that would still not justify a status quo that, on one hand, mostly forbids consensual and only moderately risky human challenge trials and, on the other hand, mostly permits harmful, lethal, nonconsensual, nontherapeutic nonhuman subjects research.

This double standard reveals inconsistency. Either we are forbidding too much human subjects research, we are permitting too much nonhuman subjects research, or we are doing some combination of the two. Of course, there is a danger in making such a point. The danger is that we might decide to resolve this inconsistency not by treating nonhumans much better but rather by treating humans much worse. I will consider that possibility below. But for now, it is enough to state that this inconsistency exists, that we need to resolve it, and that we should resolve it (at least in part) by treating nonhumans much better, not by treating humans much worse. The only question is how to flesh out the details, both in our development of new ethical standards for nonhuman subjects research and in our application of these standards.

This double standard also creates pseudo-problems. For instance, many people are currently developing human-nonhuman chimeras for research and transplantation. The basic idea is that we want animals who are human-like enough for research to be human-relevant and for transplantations to be human-compatible, yet who are nonhuman-like enough that we can permissibly harm and kill them in ways that we could never permissibly do with humans. Many people are also concerned about new moral questions that these animals raise, such as: Which moral framework should we apply to human-nonhuman chimeras? Should we apply human standards to them and treat them very well? Should we apply nonhuman standards to them and treat them badly? Or should we strike a balance between these extremes?¹¹

However, this entire line of reasoning presupposes our current, unacceptable double standard between human and nonhuman research. After all, if we eliminated this double

¹¹ For discussion of these issues, see Hyun 2003.

standard, then we would eliminate the rationale for most (in vivo) chimera use, since we would accept a presumption against harmful, lethal, nonconsensual, nontherapeutic research for humans and nonhumans alike, rather than accepting this presumption much more for humans than for nonhumans. We would also eliminate the new questions that human-nonhuman chimera use raises, since our aspirations for human-nonhuman chimeras would match our aspirations for all animals: to treat them with respect, compassion, and justice, and to replace, reduce, and refine our use of them as much as possible with those principles in mind.

6. Toward an Integrated Moral Framework

It would take much more space than I have in this chapter to develop and defend an integrated moral framework for human and nonhuman subjects research. So I will not attempt to do that here. Instead, I will attempt to lay the groundwork for this project, by describing and motivating some general features that I expect this moral framework to have, as well as some hard questions that we will need to answer as we develop it. In particular, I expect that this framework will combine the human and nonhuman paradigms by aspiring to treat all sentient animals with respect, compassion, and justice while replacing, reducing, and refining our use of them where possible. I also expect that this framework will allow us to treat humans and nonhumans differently insofar as their individual circumstances warrant that.

First, I expect that an integrated framework will aspire to treat all sentient animals with respect. To the degree that animals are capable of consent (which, at present, might apply only within humanity), that means allowing them to provide informed consent to harmful or risky research. To the degree that they are not, that means (a) adopting a presumption against harming or killing them for non-therapeutic reasons, (b) appointing a representative to make decisions on their behalf with their interests and needs in mind, and (c) allowing them to assent or dissent to particular forms of treatment to the degree that they are able. This is how we treat humans, including humans who are capable of consent and humans who are not. If we value both respect and consistency, then we should treat nonhumans similarly.

Second, I expect that an integrated framework will aspire to treat all sentient animals with compassion. That means considering animal welfare in harm-benefit analyses, as I will discuss in a moment, and it also means reducing harms and increasing benefits for nonhuman research

subjects. As with humans, we should set a limit on how much we harm nonhuman research subjects, even if they assent to the harmful activity, and even if the harmful activity is necessary for the science. We should also create the conditions necessary for nonhuman research subjects to live full, happy, and healthy lives, both during research, via species-appropriate enrichment, and after research, via species-appropriate retirement. And we should allow enough flexibility in research practices that participation can fit into a full nonhuman life.

Third, I expect that an integrated framework will aspire to treat all sentient animals with justice. In general, we should aspire to distribute the benefits and burdens of research equitably within and across species. That means, first, that we should compensate nonhuman animals for participation in research by benefiting them at least as much as (if not much more than) we harm them. It also means that we should generally avoid distributing the burdens of research disproportionately to nonhumans and the benefits of research disproportionately to humans. These principles imply that we should avoid harming nonhumans more than we can benefit them in research, which, in turn, implies that we should harm them *much* less and benefit them *much* more than we currently are, as a matter of respect, compassion, and justice.

Fourth, and relatedly, I expect that an integrated framework will include more rigorous harm-benefit analyses. In general, we should consider all relevant expected impacts, taking into account both the probability and level of benefit and harm for everyone involved. We should also take into account all relevant counterfactuals and aim to do as much good and as little harm as ethically possible rather than merely more good than harm. While the results will naturally vary from case to case, the general result is likely to be that we will permit fewer harmful nonhuman studies, since we will discover that these studies produce fewer expected benefits (once we consider counterfactual impacts) and more expected harms (once we consider long-term risks), and we will likely also set a higher bar for acceptable harm.

Fifth, I expect that an integrated framework will include more than harm-benefit analysis. For instance, even if we accept a welfarist moral theory such as utilitarianism, we might still think that rules, rights, virtues, relationships, and other such factors have an important role to play in promoting human and nonhuman welfare. For instance, when we implement systems of rules and rights, we increase the chance that humans will treat nonhuman populations well rather than use biased harm-benefit analyses to rationalize harming nonhumans to benefit humans. And when we cultivate antispeciesist beliefs, values, and habits and build antispeciesist social and

professional environments, then we create the conditions necessary for people to be motivated to follow these rules and respect these rights (John and Sebo 2020).

Sixth, I expect that an integrated framework will include the Three Rs, with much more emphasis on all three, particularly replacement. If we aspire to treat humans and nonhumans with respect, compassion, and justice, and if we aspire to both improve and supplement our use of harm-benefit analysis accordingly, then it follows that we should aspire to replace, reduce, and refine harmful human and nonhuman subjects research as much as possible. In particular, there is simply no way that we can follow the above principles while maintaining anything like current levels of harmful nonhuman subjects research. Instead, the only way that we can follow these principles is by developing animal-free alternatives as much as possible, and changing our cultural and institutional structures to accommodate them (Herrmann et al 2019).

Seventh, I expect that an integrated framework will still allow for different standards of treatment for humans and nonhumans. For instance, we might think that we can permissibly prioritize an individual human over an individual mouse, on the grounds that the human has more and stronger interests than the mouse. But even if we accept that, note two caveats. First, we might not always be permitted to prioritize humans for such reasons, since humans might not always have more, or stronger, interests than nonhumans (consider elephants, for instance). Second, even when we *are* permitted to prioritize humans for these reasons, we might still be required to prioritize nonhumans much more than we do, both individually and, especially, collectively (given how many nonhumans there are).

Eighth, I expect that an integrated framework will require us to invest in the structural conditions necessary for effective implementation. As noted a moment ago, improving our treatment of animals requires more than simply aspiring to do so. It also requires creating the cultural and institutional environments that allow us to live up to that aspiration. That means improving our education system to include less content on animal research and training with animal models and more content on animal-free alternatives and training with animal-free models – as well as more content on animal health, welfare, and rights in general. It also means creating more employment opportunities in animal-free alternatives, and creating the facilities, equipment, and technologies necessary for that to happen.¹²

¹² For more on alternatives to animal use in education, see Van Der Valk et al 1999.

Of course, fleshing out the details requires answering many extra difficult questions. First, we need to answer difficult questions about welfare. For instance, how can we estimate how much welfare different animals can have at a time and over time? Some people are exploring the idea of treating neuron counts as a proxy for welfare at a time and lifespans as a proxy for welfare over time. In that case, we might estimate that a typical human life contains about 50,000 times more welfare than a typical mouse life, and so we might assign a typical human life about 50,000 times more weight than a typical mouse life in harm-benefit analyses.¹³ But of course, it is far from clear that these are the correct proxies for welfare at a time or over time, and a lot depends on which proxies we select and why.¹⁴

We also need to answer difficult questions about rights. For example, should we treat rights as *constraints*, such that we should avoid infringing them no matter what? Or should we instead treat rights as *presumptions*, such that we should avoid infringing them ordinarily but can permissibly infringe them when the stakes are sufficiently high? Either way, we would need to forbid many nonhuman studies that we currently permit. But we might need to forbid a wider range of studies if we treat rights as constraints than if we treat them as presumptions. This is particularly true if we think that the strength of these presumptions depends on the strength of our interests, since, in that case, the bar for infringing the rights of some animals could be much higher than the bar for infringing the rights of others (Kagan 2019, Sebo 2022).

We need to answer difficult questions about many other issues as well. For instance, some people think that we can permissibly prioritize human interests because we have special duties of assistance to fellow humans, given the special relationships that we have with members of our own species (see, e.g., Brody 2012). But even if this is true, there might be a limit to how much we can permissibly prioritize human welfare, and there might also be a limit to what we can permissibly do to nonhumans in order to promote human welfare. We should also keep in mind that the reverse might sometimes be true as well; that is, we might sometimes have special duties of assistance to nonhuman animals as well, given how much we harm them, how much we benefit from them, and how much better off many of us are than many of them.¹⁵

¹³ This back-of-the-envelope estimate is based on the assumptions that a typical human has about 86 billion neurons and can live for about 79 years, whereas a typical mouse has about 70 million neurons and can live for about 2 years. But these assumptions should be questioned as well, particularly questions about nonhuman lifespans.

¹⁴ For more on cross-species welfare comparisons, see Budolfson and Spears 2020, Schukraft 2020, and Višak 2017.

¹⁵ For general discussion of these principles, see Shue 1999. For arguments that humans and nonhumans can and do have morally relevant relationships, see Gruen 2005, Palmer 2010, and Sebo 2022.

7. Assessing this Integrated Moral Framework

While it would be difficult to assess this integrated moral framework for human and nonhuman subjects research without fleshing out the details more, we can make some general observations about the benefits and risks of this approach. On one hand, this approach would aim to treat everyone as they deserve while still accommodating morally relevant differences across species. It would also allow us to improve our assessments of research, and to avoid double standards and pseudo-problems. On the other hand, there is also a risk that this approach would be simplistic and reductive, and would compromise scientific progress. I think that the benefits of a more integrated framework clearly outweigh the risks, particularly if we take care to mitigate the risks, but we will need to consider them all carefully.

Consider first the benefits of an integrated moral framework. First, this approach would aim to treat everyone as they deserve. In particular, it would require us to treat each and every research subject as an individual with rights, welfare, and morally significant relationships. Thus, it would require us to extend respect, compassion, and justice to humans and nonhumans alike, and to replace, reduce, and refine our use of them as much as possible, as a means to this end. Granted, we might still think that harming research subjects can be permissible in some cases, depending on the details of the situation and the details of our moral framework. But we would at the very least think that we should harm many fewer nonhuman animals, harm them much less, and help them much more than we currently do.

Second, this approach would still accommodate morally relevant differences across species. As I have emphasized, we can fully consider the interests and rights of, say, humans and mice while still thinking that the content and strength of our interests and rights differ dramatically – particularly if we think that rights are presumptions, that the strength of our rights depends on the strength of our interests, and that some animals can have stronger interests than others. Thus, even if we accept a unified moral framework for our interactions with, say, humans and mice in principle, we can still accept different moral frameworks for our interactions with them in practice (as we do with, say, human adults and children), provided that these different moral frameworks follow from our unified moral framework together with the facts.

Third, this approach would allow us to avoid double standards and pseudo-problems. When we think about human and nonhuman subjects research holistically, assessing each in light of the other, we are more likely to achieve consistency, since we are more likely to identify our rationalizations of nonhuman subjects research *as* rationalizations. As a result, we will not need to ask whether to apply heavily restrictive “human” standards or heavily permissive “nonhuman” standards to our interactions with human-nonhuman chimeras, since we will instead simply apply the same standards to them as to everyone: treat them with respect, compassion, and justice, and, so, replace, reduce, and refine our use of them where possible – which, in this case, means stopping this research before we start (Sebo and Parent, forthcoming).

Now consider some risks of an integrated moral framework. First, there is a risk that this approach would lead us to dehumanize humans. We already have a general tendency to dehumanize humans by treating them badly, and then rationalizing this behavior by comparing them with nonhuman animals (where the idea is that we can permissibly treat nonhumans badly, and so we can permissibly treat humans who resemble nonhumans badly as well). Humans use this rationalization to support racism, sexism, ableism, classism, and other human oppressions. In light of this tendency, we might wonder if creating a more integrated moral framework for human and nonhuman research will erode our current, fragile sense of human dignity and create a permission structure for treating humans “like animals.”¹⁶

Second, and relatedly, there is a risk that this approach would lead us to “humanize” nonhumans. We already have a general tendency to anthropomorphize nonhumans by attributing human characteristics to them whether or not they have those characteristics. We appear to develop this tendency at an early age, and we apply it to a variety of nonhumans, including not only animals but also gods. In light of this tendency, we might wonder if creating an integrated moral framework for human and nonhuman subjects research will erode our current appreciation of the many morally relevant differences across species, with the result that we attribute human interests to nonhuman research subjects much more than we should, and attribute (distinctively) nonhuman interests to them much less than we should.¹⁷

¹⁶ For more on dehumanization, see Smith 2020. For more on dehumanization and speciesism as they relate to racism, sexism, and ableism, see, respectively, Ko and Ko 2017, Adams 1990, and Taylor 2017.

¹⁷ For general discussion of anthropomorphism, see should Daston and Mitman 2005.

Third, there is a risk that this approach would compromise scientific and medical progress. Our current systems of science and medicine are based on massive amounts of harmful, lethal, nonconsensual, nontherapeutic nonhuman subjects research. As noted above, we currently see nonhuman subjects research as part of the best, if not the only, path to approval for many foods and drugs, and while animal-free alternatives are available in some cases, they might not be available in all cases. Thus, if we held all research to high standards, then we might slow scientific and medical progress (under current regulations) or compromise the safety and efficacy of new products (under new regulations). Either way we would be replacing one set of risks and harms with another, likely burdening nonhumans less and humans more.

My own view is that the benefits of an integrated moral framework decisively outweigh the risks. The benefits of an integrated moral framework are difficult to overstate. The current research ethics paradigm is neither scientifically nor ethically optimal, given how different humans and nonhumans are and how much nonhumans suffer in research. Yet we continue with it anyway because of cultural and institutional bias, ignorance, and path dependence. If we improve and integrate oversight of human and nonhuman subjects research in the ways that we have discussed here, then we can still make progress in science and medicine while reducing harm to nonhumans via research, reducing risks for humans via false positives and negatives, and reducing risks for nonhumans via reinforced speciesism.

Meanwhile, we can mitigate the risks. Consider each in turn. First, we can mitigate the risk that an integrated moral framework will lead us to dehumanize humans. As we have seen, we can accept a unified framework for humans and nonhumans in principle while still accepting different frameworks for them in practice, since we might think that our rights can vary with our interests and relationships. And if we emphasize that our goal is to extend dignity to nonhumans rather than restrict it from humans, then we can mitigate the risk that this extension will erode our sense of human dignity. Granted, it might erode our sense of human *supremacy*, but that would be appropriate. When the status quo involves massive and unnecessary exploitation and extermination of vulnerable others, the status quo needs to change.

Second, and relatedly, we can mitigate the risk that an integrated moral framework will lead us to “humanize” nonhumans. Again, we can accept a unified framework for humans and nonhumans in principle while still accepting different frameworks for them in practice. And if we make sure to do this work together with research in cognitive ethology and comparative

psychology, taking care to note the similarities as well as the differences across species, then we can mitigate the risk that we will treat nonhumans as more human-like than they are. Granted, we might treat nonhumans as more human-like than we currently do, but, again, that would be appropriate. When the status quo involves excessive *anthropodenial*, at least some additional anthropomorphism might be necessary to establish an equilibrium.

Third, we can mitigate the risk that an integrated moral framework will compromise scientific and medical progress. We already have alternatives to many current research methods, and we will likely be able to develop more with time (Herrmann et al 2019). Granted, there might be some trade-offs between ethics and science during the transition, and we should take these trade-offs seriously. But while taking these trade-offs seriously might require harming nonhumans for the greater good in some cases (particularly if we treat rights as presumptions rather than constraints), it might also require not doing so in other cases. And if we are prepared to delay some scientific advances for the sake of human rights, welfare, and justice, then we should be prepared to do the same for the sake of nonhuman rights, welfare, and justice.

In short, if we build an integrated moral framework for human and nonhuman subjects research with the general principles that I have described here in mind, then we can strike a much better balance between integration and fragmentation in research ethics. Our treatment of humans and nonhumans can be integrated in that they can flow from a unified set of considerations regarding respect, compassion, justice, replacement, reduction, and refinement. And they can be fragmented in that they can still have different implications for different individuals, depending on, for instance, the nature of their interests, rights, and relationships. The upshot might be priority for humans in some respects, but not nearly as much as we currently enjoy. This is a change that we should welcome rather than continue to resist.¹⁸

References

Carol J. Adams, *The Sexual Politics of Meat: A Feminist-Vegetarian Critical Theory* (New York, NY: Bloomsbury Academic, 1990).

¹⁸ Thanks to Carolyn Neuhaus and Brendan Parent for helpful feedback on the penultimate draft of this chapter, and thanks to Erick Valdés and Juan Alberto Lecaros for all their hard and great work editing this chapter and book.

Robert Bass, "Lives in the Balance: Utilitarianism and Animal Research," in *The Ethics of Animal Research: Exploring the Controversy*, ed. Jeremy R. Garrett (Cambridge, MA: MIT Press, 2012), pp. 81-106.

Bracken-Roche D, Bell E, Macdonald ME, Racine E. "The concept of 'vulnerability' in research ethics: an in-depth analysis of policies and guidelines," [published correction appears in *Health Res Policy Syst*. 15(1):29 (2017)]. *Health Res Policy Syst*;15(1):8 (2017).

Baruch A Brody, "Defending Animal Research: An International Perspective," in *The Ethics of Animal Research: Exploring the Controversy*, ed. Jeremy R Garrett (Cambridge, MA: MIT Press, 2012), pp. 53-66.

Mark Budolfson and Dean Spears, "Public Policy, Consequentialism, the Environment, and Non-Human Animals," in *The Oxford Handbook of Consequentialism*, ed. Douglas W. Portmore (New York, NY: Oxford University Press, 2020), pp. 592-615.

Ewen Callaway, "Dozens to be deliberately infected with coronavirus in UK 'human challenge' trials, *Nature* 2020: <https://www.nature.com/articles/d41586-020-02821-4>.

Lorraine Daston and Gregg Mitman, *Thinking with Animals* (New York: Columbia University Press, 2005).

Matthias Eggel, Carolyn Neuhaus, and Herwig Grimm, "Reevaluating benefits in the moral justification of animal research: A comment on 'Necessary conditions for morally responsible animal research,'" *Cambridge Quarterly in Healthcare Ethics* 29:1 (2020): 131-143.

Arianna Ferrari, "Contesting animal experiments through ethics and epistemology: In defense of a political critique of animal experimentation," in *Animal Experimentation: Working towards a paradigm change*, ed. Kathrin Herrmann and Kimberley Jayne (Boston: Brill, 2019), 194-206.

Samuel Fleischacker, *Divine Teaching and the Way of the World* (Oxford: Oxford University

Press, 2011).

Christine Grady, "Payment of clinical research subjects," *The Journal of Clinical Investigation* 115:7 (2005): 1681-1687.

Lori Gruen, *Entangled Empathy* (New York: Lantern Books, 2005).

Kathrin Herrmann, Francesca Pistollato, and Martin L Stephens, "Beyond the 3Rs: Expanding the use of human-relevant replacement methods in biomedical research." *ALTEX*. 36:3 (2019): 343-352.

Insoo Hyun, "Illusory Fears Must Not Stifle Chimaera Research," *Nature* 537, 281 (2016).

Tyler John and Jeff Sebo, "Consequentialism and Nonhuman Animals," in Douglas Portmore, ed., *The Oxford Handbook of Consequentialism* (Oxford: Oxford University Press, 2020), pp. 564-591.

Shelly Kagan, *How to Count Animals, More or Less* (Oxford, UK: Oxford University Press, 2019).

Holly Kantin and David Wendler, "Is there a role for assent or dissent in animal research?" *Cambridge Quarterly of Healthcare Ethics* 24:4 (2015): 459-72.

Aph Ko and Syl Ko, *Aphro-Isms: Essays on Pop Culture, Feminism, and Black Veganism from Two Sisters* (Brooklyn, NY: Lantern Books, 2017).

Clare Palmer, *Animal Ethics in Context* (New York, NY: Columbia University Press, 2010).

John Rawls, "The idea of an overlapping consensus," *Oxford Journal of Legal Studies* 7:1 (1987), 1-25.

Tom Regan, *The Case for Animal Rights* (Berkeley, CA: University of California Press, 1983).

David Resnik, *The Ethics of Research with Human Subjects* (New York: Springer, 2018).

Jason Scott Robert and Françoise Baylis, "Crossing species boundaries," *American Journal of Bioethics* 3.3, 1-13 (2003).

W. M. Russell, R. I. Burch, and C. w. Hume, *Principles of Humane Experimental Technique*. (Hertfordshire: Universities Federation for Animal Welfare, 1959).

Jason Schukraft, "How to Measure Capacity for Welfare and Moral Status," *Rethink Priorities*, June 1, 2020.

Jeff Sebo, *Saving Animals, Saving Ourselves* (Oxford University Press, 2022).

Jeff Sebo and Brendan Parent, "Human, nonhuman, and chimera research ethics," [publication information forthcoming].

Jeff Sebo and David DeGrazia, "Can knowledge itself justify harmful research?," *Cambridge Quarterly in Healthcare Ethics* 29:2 (2020): 302-307.

Henry Shue, "Global Environment and International Inequality." *International Affairs* 75 (1999): 531–545.

Peter Singer, *Animal Liberation: A New Ethics for Our Treatment of Animals* (New York, NY: Avon, 1975).

David Livingstone Smith, *On Inhumanity* (Oxford: Oxford University Press, 2020).

Katy Taylor and Laura Rego Alvarez, "An estimate of the number of animals used for scientific purposes worldwide in 2015," *Alternatives to Laboratory Animals* (2020).

Sunaura Taylor, *Beasts of Burden: Animal and Disability Liberation* (New York: The New Press, 2017).

Jan Van Der Valk et al., “Alternatives to the Use of Animals in Higher Education: The Report and Recommendations of ECVAM Workshop 33,” *Alternatives to Laboratory Animals* 27, no. 1 (1999): pp. 39-52.

Tatjana Višak. “Cross-Species Comparisons of Welfare,” in *Ethical and Political Approaches to Nonhuman Animal Issues*, A. Woodhall, G. Garmendia da Trindade (Springer International Publishing, 2017).