Utilitas (2026), **38**, 143–154
doi:10.1017/S0953820826100272

CAMBRIDGE UNIVERSITY PRESS

ARTICLES

Harm-Benefit Analysis for Animal Experiments Is Not Utilitarian

Nico Dario Müller[1] and Tristan Katz[2]

[1]Department of Media, Art, Philosophy, University of Basel, Switzerland and [2]Environmental Sciences and Humanities Institute, University of Fribourg, Switzerland
Corresponding author: Nico Dario Müller; Email: nicodario.mueller@unibas.ch

(Received 11 February 2025; revised 8 December 2025; accepted 14 January 2026; first published online 25 March 2026)

Abstract
The claim is commonplace that harm-benefit analysis (HBA), a weighing procedure widely used in ethics reviews of animal experiments, is utilitarian. We argue this is false and misleading for three reasons: (1) HBA does not compare, let alone maximize, utility across different options, but merely assesses whether the consequences of one option are net-positive, thereby ignoring opportunity costs; (2) HBA does not aggregate utility coherently, as it allows for varying degrees of speculation in the assessment of harms and benefits; (3) HBA is not concerned with moral evaluation or moral goodness. From our discussion, we derive positive suggestions for how to improve animal experimentation policy and public communications about it. Most straightforwardly, scholars and institutions should stop claiming that HBA is "utilitarian."

Keywords: Harm-benefit analysis; animal experimentation; animal research; regulation; utilitarianism

1 Introduction

It is commonly claimed that harm-benefit analysis (HBA), a weighing procedure widely used in ethics reviews of animal experiments, is "utilitarian" (for examples, see section 2). This conveys a degree of philosophical respectability. Controversies notwithstanding, utilitarianism is one of history's most influential schools of ethical thought. We argue that HBA does not merit this pedigree. The claim that it is a "utilitarian" procedure is false, misleading, and perpetuates misconceptions about both HBA and utilitarianism. Considering exactly why the claim is wrongheaded, however, helps to see how HBA can be improved, along with the regulation of animal experimentation as a whole. Our discussion will proceed as follows. First, we will introduce HBA and discuss examples of it being called "utilitarian" in the literature (section 2). Second, we will endorse a broad definition of "utilitarianism" (section 3) and argue that HBA is not utilitarian for at least three key reasons: It ignores opportunity costs (section 3.1), it lacks a coherent aggregation rule (section 3.2), and it is not concerned with moral goodness (section 3.3). We will offer some suggestions for how animal experimentation regulation

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

and communication about it could be improved based on our discussion (section 4). Section 5 concludes.

2 The claim that HBA is utilitarian

HBA first became part of animal testing regulations in the United Kingdom with the Animals (Scientific Procedures) Act of 1986 (Brønstad et al. 2016, p. 11). Today, many countries' regulations and influential guidelines require some form of HBA before an animal experiment can be licensed (Brønstad et al. 2016; Laber et al. 2016; Bayne and others 2024).

Although more fine-grained HBA procedures exist (Brønstad et al. 2016; Laber et al. 2016), an HBA procedure consists of three basic parts. In the first part, the responsible committee assesses the harm to be inflicted on the animals, often using a classification system for degrees of severity (Brønstad et al. 2016, p. 7). It identifies modulating factors – that is, respects in which the study might be changed to inflict less harm on animals (Laber et al. 2016, p. 23) – in order to assess the applicant's implementation of the "three Rs" of "replace, reduce, refine" (Russell and Burch 1959). Thus, the applicant must explain why alternatives to animal-reliant methods are not adequate for their research endeavor, how they plan to minimize the number of animals used, and what measures they will take to decrease the harm inflicted. In the second part, the committee considers the prospective benefits of the study. The applicant typically explains the contribution of the study as either an advancement in basic research or the solution of a societal problem (e.g., finding a safe and effective treatment for some disease). Again, modulating factors are identified, such as how objectives will be met and how well the study fits into an existing body of research (Laber et al. 2016, p. 27). At this step, committees often devote great attention to scientific quality, guided by the idea that poorly designed studies are unlikely to yield benefits that could justify harm to animals – indeed, discussions of scientific potential and quality often dominate committee discussions (Olsson et al. 2015, p. 35). In the third and final part of HBA, committee members deliberate about whether the benefits outweigh the harms. Although technical frameworks like Bateson's cube (Bateson 1986) can aid them in this process, there is no straightforward "algorithm" that could dictate decisions, in part because harms and benefits do not admit of a common unit and cannot literally be weighed (Laber et al. 2016, p. 30). Rather, each member forms their own opinion and the committee's verdict is determined by majority vote, though consensus would be desirable (ibid.).

The claim that it is "utilitarian" to evaluate animal experiments by weighing up their harms and benefits has become increasingly prominent over the last half century. Before, the word "utilitarianism" was sometimes used more loosely in discussions of animal experimentation: in the sense of "ruthlessness" or "unscrupulousness." For example, C. S. Lewis wrote that "the victory of vivisection marks a great advance in the triumph of ruthless, non-moral utilitarianism over the old world of ethical law" (1947, p. 11). Charles W. Hume, who oversaw the project from which Russell and Burch's (1959) "three Rs" concept would emerge, primarily associated utilitarianism with the idea that the ends justify the means. He pointed out that British law prohibited torturing humans irrespective of its potential benefits, as well as animal experiments above a certain upper limit of pain, and thus was not fully utilitarian (1951, pp. 88–89). Although Hume did refer to utilitarianism as a philosophical principle ("[. . .] which prescribes the greatest good of the greatest number as the sole ethical objective" 1951, p. 88), he

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

largely used it as a shorthand for the unscrupulous willingness to do whatever brings more benefit than harm (see also Hume 1962, p. 129).

The philosophical tradition of utilitarianism became more central to debates about animal experimentation in the wake of Peter Singer's "Animal Liberation" (1975). Around that time, moral debates about animal experimentation intensified in general (Phillips and Sechzer 1989, p. 43) and philosophers became more involved (ibid., p. 63). Singer's view was that a utilitarian should be critical of almost all animal experimentation because it did not lead to the greatest utility for all affected (Singer 1975, Ch. 2). A debate ensued when other philosophers, most notably Carl Cohen (1986, 2001), defended animal experimentation on utilitarian grounds.

At a 1980 meeting with various scientists and philosophers, Patrick Bateson – inventor of the aforementioned Bateson's cube (Bateson 1986), a framework often used in HBA (Brønstad et al. 2016, pp. 11–12) – is reported to have said that "the 'good' scientist adopts a more-or-less utilitarian view, and allows a little more suffering to very high quality research" (Cherfas 1980, p. 1003). Some have understood him to appeal to utilitarianism in the sense of the philosophical tradition (Niemi 2019), though he might have used the term more loosely, like Lewis and Hume.

In any case, since the late 1980s, a fairly steady stream of literature about animal experimentation can be traced in which the philosophy of utilitarianism is equated with the balancing of harms against benefits (e.g., Sumner 1988, p. 162; Orlans 1993, p. 121; Beauchamp et al. 1998, p. 79; Würbel 2009, p. 119; Graham and Prescott 2015, p. 20; Taylor 2018, p. 148; Beauchamp and DeGrazia 2020, p. 30). The more specific claim that HBA is a utilitarian procedure is virtually ubiquitous in literature about HBA (Donnelley and Nolan 1990, p. 5; Knight 2011, p. 180; Bout et al. 2014, p. 411; Griffin et al. 2014, p. 266; Graham and Prescott 2015, p. 21; Walker 2016, p. 29; Davies et al. 2017, p. 15; Davies 2018, p. 57; Niemi 2019, p. 341; Gutfreund 2020, p. 2; Pollo and Vitale 2020, p. 200; Mertz et al. 2024, p. 11).

The claim that HBA is utilitarian does not usually carry much explicit argumentative weight, but it exerts framing power. Its presence confers the respectability of a centuries-old philosophical tradition upon HBA and suggests that the framework is ethically justified, at least according to one highly influential way of looking at morality. This, we will now argue, is undeserved.

3 Why HBA is not utilitarian

In the following, we rely on a broad notion of utilitarianism (following Scarre 1996). Utilitarianism is an approach to normative ethics that is welfarist, consequentialist, aggregative, and comparative (maximizing, satisficing, or scalar). It is *welfarist* in that it evaluates acts or rules by how they affect welfare, i.e., how well lives are going (Scarre 1996, pp. 4–9). It is *consequentialist* in that it evaluates acts or rules by the goodness of their consequences (Scarre 1996, pp. 10–14). It is *aggregative* in that it judges acts or rules by the sum goodness or badness of their overall consequences for the welfare of all affected, though this scope might sometimes be restricted, e.g., to humans only (Scarre 1996, pp. 14–18). Last but not least, utilitarianism is *comparative* in that it evaluates acts or rules by how much utility they bring about, relative to other available courses of action. Classical utilitarianism does so by selecting for acts or rules that bring about the greatest aggregate utility possible (maximizing utilitarianism). Less demandingly, utilitarianism can select for consequences that are good enough by some appropriate standard (satisficing utilitarianism), even if they might fall short of achieving

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

maximum utility (Scarre 1996, pp. 18–23; Slote and Pettit 1984). Utilitarianism can also grade acts or rules on a scale from bad to good, depending on how they compare to other available options, doing away with notions of permissibility and obligation (scalar utilitarianism, Norcross 2006). Within the four defining characteristics, there is leeway for different views, e.g., on which intrinsic goods a utilitarian considers relevant for welfare and how she aggregates consequences.

There are, conversely, some views that do not define utilitarianism, although they are often associated with it. First, utilitarianism need not define the good in terms of pleasure and pain, although the classical utilitarians did. Others, particularly the "ideal utilitarians" Hastings Rashdall and G. E. Moore, recognized other intrinsic goods, such as beauty (Driver 2022, Ch. 4; Scarre 1996, Ch. 5.2). Utilitarianism by our definition thus does not have to be hedonistic or value-monist.

Second, utilitarians need not think that aggregation is linear, in the sense that adding the same amount of a good thing to any set of consequences raises the overall goodness of that set by the same amount. G. E. Moore put forward a theory of "organic wholes" according to which a good's overall goodness can differ from the sum goodness of its constituents (Moore 1922, §§ 18–22). His example is that an experience of a beautiful object may be very good indeed, even though its constituents – a conscious experience and the existence of a beautiful object – are both barely good at all when considered separately.

Third, utilitarians need not understand goods as being "homogeneous" in two respects. (a) They need not think that all pleasures are equally good. John Stuart Mill famously thought there were "higher" and "lower" pleasures of unequal goodness (Mill 2015, pp. 121–125). (b) Utilitarians need not think that the same pleasure or pain is equally good irrespective of who has it (as endorsed by Singer 1975, in the form of the principle of equal consideration of interests). With an "organic wholes" theory, a utilitarian could argue that the same pleasure in a different subject results in greater or lesser overall goodness (e.g., in someone who deserves it versus someone who does not).

Our notion of utilitarianism is specific enough to set utilitarianism apart from other approaches to moral philosophy, such as egoism, Kantianism, and virtue ethics. An egoist would disagree with the aggregative aspect of utilitarianism, considering only the agent's own welfare. A Kantian or virtue ethicist would disagree with consequentialism because they would not use consequences as the yardstick of moral evaluation, but rather the presence of a good will or the virtues (though consequences may still play an important role in these approaches). On the other hand, our notion of utilitarianism is broad enough to capture many varieties of the approach, not just the classical ones. It specifically includes non-maximizing and non-hedonistic variants that some might consider edge cases of the theory (Scarre 1996, p. 114). This is important for present purposes because our aim is not just to show that HBA is misaligned with specific varieties of utilitarianism, but that it differs so fundamentally from utilitarian thinking that it is false and misleading to claim that the former is based on the latter.

3.1 HBA ignores opportunity costs

The first and perhaps most important divergence of HBA from utilitarian thinking is that it ignores opportunity costs. Although utilitarianism is often understood as aiming merely at net-positive aggregate consequences, this is not correct. Utilitarians aim at aggregate consequences that compare well against other available alternatives. Typically, they do so by aiming either to maximize utility or to satisfice utility using an

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

appropriately demanding standard. Both require looking not just at whether a proposed act or rule has net-positive aggregate consequences, but also at how these consequences compare against the best available alternative. To maximize utility, an agent must choose, considering all possible acts or rules, the one that has the best aggregate consequences. To satisfice utility, it suffices that the agent searches until she finds an option that is good enough; still, a range of candidate acts or rules must be considered. A scalar utilitarian, meanwhile, would focus directly on comparisons between available options.

In the case of an animal experiment, the utilitarian question is therefore *not* whether its benefits outweigh its harms, but how its consequences compare to those of other available options (see Sapontzis 1988, p. 190). This would require thinking about the best possible thing one could do with the resources the study proposes to use, such as money, time, and person-power, as Bass (2012, p. 88) has pointed out. This would mean asking "is this animal study the best possible study we can conduct with these resources?" But it would also mean asking more broadly whether our resources are best spent on research at all, given that we might be able to achieve greater benefits by spending them on implementing solutions we already know (ibid.). These are not trivial questions, especially in a world where the research agenda often diverges from societal needs and the benefits that flow from research are generally unfairly distributed (Kitcher 2001; Reiss and Kitcher 2009).

Contrary to these utilitarian requirements, HBA does not compare different future courses of action. It compares the consequences of one future course of action – conducting the animal experiment – against the status quo; that is to say, against net-zero utility. In fact, just this aspect of animal experimentation regulation has been criticized: "The choice is not between [animal study] X and nothing, but between X and myriad other things" (Lauwereyns 2018, p. 107). To be clear, however, HBA does not compare one course of action (the animal study) to another (doing nothing). After all, the course of action described as "doing nothing" – having the scientists idle in their labs, presumably – could itself have net-positive or net-negative consequences, but this is not assessed in HBA. The truth is that HBA does not compare different courses of action at all, but only checks whether the consequences of one particular option are net-positive.

One might wonder whether HBA is based on an especially undemanding version of satisficing utilitarianism, where anything above net-zero utility counts as good enough. But no serious moral theory can approve of just any course of action as long as its consequences are net-positive. Such a principle would fail to meaningfully guide moral agents. For instance, it could not tell an agent whether to donate one dollar, a hundred dollars, or ten million dollars to a worthy cause, even if the agent is a billionaire. Similarly, it would not tell an agent, faced with some great disaster, whether to help only one person in need, to help several, or as many as they can manage, even if they had desperately needed (e.g., medical) expertise. All options are net-positive, after all. Not only is it implausible that all these options should be *equally good*, but the less helpful options seem to be *bad* despite having net-positive consequences. Thus, even if we were to widen our definition of utilitarianism so that selecting for net-positive consequences became a conceivable variety of it, this would be a patently implausible variety that nobody, to our knowledge, actually endorses.

HBA neither maximizes nor satisfices utility to an appropriate standard, nor does it rank animal experiments by comparing their utility to that of alternative options. Contrary to this, the HBA literature has made claims such as "HBA is a form of decision-

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

making that uses moral reasoning based on utilitarianism; i.e. aiming for the maximum balance of benefits over harms for all affected" (Davies et al. 2017, p. 15; similarly Davies 2018, p. 57) and "HBA is based on the assumptions of maximizing utility for the majority where human interests count most and is an essential part of the ethical review" (Brønstad et al. 2016, p. 17). It is very hard to see how a decision procedure could be "based on" a principle of utility maximization if it does not even consider other available courses of action.

The (false) claim that HBA selects for maximum utility can easily be conflated with another (true) claim, that HBA aims to maximize the utility of one given experiment (e.g., Griffin et al. 2014, p. 265; Balls 2021, p. 184). By checking whether a study's harms can be further mitigated and its scientific quality improved, HBA does take steps to this effect. But this does not entail that the procedure selects for maximum utility, just like pushing one runner to their personal best does not entail having found the fastest person on Earth (or a fast-enough person by some relevant standard). No serious moral theory could approve of just any experiment as long as its utility is maximized within its own limits, even if the condition was added that aggregate consequences must be net-positive. The problem would remain that the theory fails to compare the action to any alternatives, making it implausible.

3.2 HBA lacks a coherent aggregation rule

Presumably, the reason why HBA is so often called "utilitarian" is that it looks like an aggregative procedure. But it is not obvious that HBA aggregates at all, and if it does, its mode of aggregation does not follow the kind of coherent rule that utilitarianism would require.

First, it might be thought that, since HBA helps the agent to deal with moral trade-offs, it must involve aggregation. But in fact, there are many non-aggregative ways of dealing with moral trade-offs, since all moral theories require a way to resolve apparent conflicts between obligations or claims. For example, Tom Regan's animal rights theory, which is explicitly anti-utilitarian, is committed to the "miniride" principle, according to which the smaller of two groups of victims should be harmed if the harms they stand to suffer are of comparable severity (Regan 2004, pp. 305–307). The classic debate sparked by John Taurek's article "Should the Numbers Count?" (1977) was populated by many non-utilitarian philosophers who argued that one should resolve at least certain moral conflicts by saving the greater number of people at stake (e.g., Kavka 1979; Sanders 1988; Kamm 1993; Scanlon 1998; Timmermann 2004). So the mere fact that HBA involves a moral trade-off is not a reason to consider it utilitarian.

Second, assuming that HBA is indeed a utility-aggregating approach, it does not aggregate utility according to any coherent rule. It treats harms and benefits unequally in unexplained ways.

This problem should be set apart from others that are similar, but not the same. One is the problem that HBA involves "comparing apples and oranges," which is well-known in HBA literature (Grimm 2015; Laber et al. 2016, p. 30; Brønstad et al. 2016, p. 3). For example, the good of knowledge is weighed against the evil of pain. This makes for a divergence from hedonistic utilitarianism, in which all utility can be broken down to a common denominator in terms of the only intrinsic good, pleasure. But there are forms of utilitarianism that recognize intrinsic goods beyond pleasure (Driver 2022, Ch. 4). Aggregation requires that intrinsic goodness is one unitary property, but several distinct things might exhibit it (see Mason 2023). However, a major challenge for these forms of

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

utilitarianism is to give guidance on how to aggregate incommensurable goods (Scarre 1996, pp. 120–121). So even though HBA is not misaligned with all forms of utilitarianism due to the "apples-vs-oranges" problem, it only aligns with varieties of utilitarianism that are obscure about just the kind of weighing HBA involves. Instead of vindicating HBA, a utilitarian pedigree would then only call it into question.

Another related issue with the way utility is aggregated in HBA is that human interests are given greater weight than animal interests (Brønstad et al. 2016, p. 16; Mertz et al. 2024, p. 11). In defense of HBA, utilitarianism is compatible with speciesism or other views that tend to give preference to (most) humans over other animals, such as Mill's theory of higher and lower pleasures (Mill 2015, pp. 121–125). Of course, one can take issue with these views, and it is worth noting that HBA as practiced aligns with the less animal-friendly varieties of utilitarianism at best. But giving greater weight to human interests, or to interests humans are more likely to have, can in principle be part of a utilitarian's aggregation rule.

Our concern is rather that HBA involves treating the very same kind of utility differentially in unexplained ways during aggregation. Specifically, while HBA is generous about speculative benefits, it only considers harms that are certain to arise. Consider first the essential, if mostly tacit role that speculation about the future plays on the benefits side. Single studies do not answer scientific questions on their own, but contribute information to a landscape of research that, in concert, can over time give increasingly definite answers (see Grimm et al. 2017). But having answers to crucial questions, e.g., about how to treat a certain disease, still does not guarantee that the societal benefit (a safe and effective treatment for the patients that need it) will manifest. This depends on an environment comprising, among other things, suitable regulations and markets. The inference from the statement that a study addresses a crucial question about a disease that affects many patients, to the claim that the study is important to conduct, tacitly commits one to a whole host of speculative assumptions about how the scientific and societal environment will react to the information produced.

Perhaps this amount of speculation is inevitable in practical decision-making. But it stands in stark contrast to the assessment of harms, in which HBA does *not* grant the same speculative freedom. HBA considers the harms directly inflicted on animals in the course of a study, but not any further harms that the study *might* help to bring about. For example, if a committee member raised the worry, without any direct evidence to support it, that a certain animal study might contribute to an overemphasis on one area of scientific activity, draining resources needed for other, neglected areas, she would be likely to earn confused looks from her colleagues. The same is true of speculative worries such as that the study might erode public trust in research institutions, that it might encourage even more harmful studies, or that it might contribute to a lock-in of science into certain types of animal experiments. These worries would be highly speculative and somewhat arbitrary. But they would not be speculative and arbitrary in a way that the assessment of potential benefits is not.

Why then treat the goodness of positive utility that might materialize differently from the badness of negative utility that might materialize? Utilitarianism does not necessarily require linear aggregation, but it does require coherent aggregation rules. A utilitarian procedure would either assess harms and benefits equally given the same degree of speculation and the same degree of cold, hard facts, or require some theory-internal justification to treat the two differently. But it is hard to see what this justification would be.

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

3.3 HBA is not about moral goodness

There is a third reason that at least makes it misleading to call HBA "utilitarian:" The procedure is not concerned with *moral* evaluation, nor should it be. Utilitarianism, as a form of consequentialism, is first and foremost about what agents morally ought to do, or what would be morally good for them to do. By contrast, the permissibility that is assessed using HBA is best understood as legal or institutional permissibility, not moral permissibility. The very idea that animal experimentation regulation should be based on moral theories conflicts with the liberal consensus view that particular conceptions of the good should not directly shape public institutions (Rawls 1999). Given the diversity of moral thought in society as well as among scholars, any particular moral perspective lacks the necessary consensus support. Tying animal research licenses to utilitarian standards of good action would be akin to tying them to Catholic Church doctrine. In short, a moral theory provides the wrong kind of justification for a procedure established by law.

The normative foundations of HBA should instead be sought in legislation and legal reasoning. The immediate justification for conducting HBA is that there is a legal basis for it, as in the UK's Animals (Scientific Procedures) Act of 1986. In some legal systems, this may be all there is to it, and HBA would cease to have any justification if the law stopped demanding it. In other systems, such as in Switzerland, HBA can be understood as a specification of a more general approach to handling conflicts between goods that are protected by the country's constitution (Gerritsen 2022). The reasoning for HBA as required under Swiss law is essentially that the state has conflicting responsibilities to both protect animals and promote scientific research. A weighing approach is customary in such cases of conflict, guided by the view that sacrificing one good for the other is optimal for the protection of both goods. Thus, HBA is justified by the fact that the constitution is committed to these goods and HBA is the designated way to resolve conflicts between them.

The claim that HBA is utilitarian also runs the risk of depoliticizing animal experimentation regulation. It either suggests that there is an ethical consensus about what animal experiments should be done or that one particular view deserves political priority over all others. Either way, it downplays the extent to which this regulation is the result of political negotiation rather than universal moral agreement (for an account of the British negotiations, see Lyons 2013). Thus, the claim that HBA is utilitarian is not only false and misleading, but can also have pernicious practical consequences.

4 Suggestions for animal experimentation policy

So far, we have argued that HBA diverges from core features of utilitarian thinking and that it is a mistake to justify the licensing procedures of public institutions with a moral theory in the first place. We do not argue that the system should be made more utilitarian, as this would merely repeat the mistake discussed in section 3.3. But the limitations of HBA we have highlighted are problematic for broader moral and institutional reasons. We thus offer three suggestions.

First, even non-utilitarian policy-makers have good reason to reform animal experimentation regulations so that opportunity costs are better taken into account. A mere weighing approach may be optimal when we are faced with one-off conflicts, e.g., when an unforeseeable natural disaster requires blocking a public road so that emergency services can use it efficiently, violating citizens' right to free movement.

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

In such cases, the best we might be able to do is weigh the goods at stake and sacrifice one for the other. But when such conflicts occur frequently and could be avoided through prospective measures (e.g., building an additional road), doing so is better than waiting until the disaster happens and then weighing the goods. Just this, it seems, is the situation in animal experimentation. So policy-makers should take prospective measures to render the goods of animal welfare and science more mutually compatible over time, specifically through strong and strategically planned support for ways of doing science that do not harm animals (see Müller 2024).

Second, speculation in HBA should be treated more consistently. It is hard to justify the differential use of speculation in the assessment of benefits and harms in HBA, from utilitarian as much as other perspectives. Adjustments are possible on both the harms and the benefits side, by increasing the latitude for speculation about the potential harms of a study or by reducing the amount of speculation in the assessment of benefits. The latter might be achieved to some extent by using the "Theory of Change" (TOC) approach to prioritize high-impact research questions (Vogel 2012). A TOC requires backwards planning from a distant goal, identifying the necessary and sufficient steps along the way (ibid., p. 1). The procedure helps to articulate tacit assumptions made about the environment of prospective research contributions which, while inevitably speculative, can be assessed for plausibility. For example, it may turn out that a research contribution is likely to bring societal benefits only if certain regulations will change or if industry picks up the study's solution to a problem, and there may exist relevant evidence as to how likely these events are. The more tenuous a claim to future benefits is, the less it should be considered to counterbalance any harms. Comparing the TOCs of different research projects could also help to assess opportunity costs at the funding stage, encouraging resources to flow towards projects with the greatest plausible utility. However, since it is exceedingly difficult to predict whether research will yield the expected benefits, especially over a timespan of longer than five years (Shaw 2022), prospective benefit assessment is bound to remain speculative for many studies even if the mechanisms behind potential benefits are more clearly articulated in a TOC. The same, we can presume, is the case for *some* harms that are contingent on various environmental factors. Thus, leveling the playing field in HBA requires allowing more speculation on the harms side.

Third, public institutions should avoid talk of HBA as a *morally* justified procedure of licensing animal studies. As we have emphasized, whether any particular moral perspective approves of HBA is irrelevant for the legitimacy of a public institution in a liberal state. Instead, institutions would do well to communicate only that HBA has a legal or institutional basis and to make transparent that different opinions about its adequacy are possible and actually being endorsed.

5 Conclusion

In this paper, we have argued that it is false and misleading to call HBA a "utilitarian" procedure for three reasons: First, utilitarianism requires the consideration of opportunity costs, but HBA ignores them: HBA is a procedure that compares the utility of one course of action against net-zero, not against the utility of alternative courses of action. Second, utilitarianism requires that aggregation follows coherent rules, but HBA treats harms and benefits unequally in unexplained ways. The assessment of benefits is deeply speculative, while the assessment of harms is restricted to cold, hard facts. Third, utilitarianism is a particular family of theories about what is morally good, but HBA is

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

not about moral goodness, and even if it were, no particular moral theory should set the standards for what is allowed in a liberal state. These three reasons lead us to the suggestions that animal experimentation policies should be reformed to consider opportunity costs; that the role of speculation in HBA should be revisited; and that public institutions should avoid portraying HBA as a morally justified procedure. The claim that HBA is "utilitarian" bestows philosophical respectability on a procedure that does not deserve it. It perpetuates misconceptions about HBA and utilitarianism. And it unduly depoliticizes the regulation of animal experimentation. In short, the claim is false, misleading, and unhelpful, and scholars should stop making it.

Financial support. Funding provided by Swiss National Science Foundation, National Research Programme 79 "Advancing 3R: Research, Animals and Society" (Grant No. 407940_214850), and Swiss National Science Foundation Project "Principles for Ethical Decision-Making in Environmental Practice" (Grant No. 197363).

Competing interests. The authors declare none.

References

Balls, Michael, 'It's Time to Include Harm to Humans in Harm-Benefit Analysis – But How to Do It, That Is the Question', *Alternatives to Laboratory Animals*, **49**.5 (2021), 182–196. https://doi.org/10.1177/02611929211062223.

Bass, Robert, 'Lives in the Balance: Utilitarianism and Animal Research', in *The Ethics of Animal Research: Exploring the Controversy*, ed. by J. R. Garrett (Cambridge, MA: MIT Press, 2012), 81–106. https://doi.org/10.7551/mitpress/9780262017060.003.0006.

Bateson, Patrick, 'When to Experiment on Animals', *New Scientist*, **109**.1496 (1986), 30–32.

Bayne, Kathryn and others, 'Legislation and Oversight of the Conduct of Research Using Animals: A Global Overview', in *The Ufaw Handbook on the Care and Management of Laboratory and Other Research Animals*, ed. by Huw Golledge and Claire Richardson (Chichester: John Wiley & Sons, 2024), 9th edn., 101–121. https://doi.org/10.1002/9781119555278.ch8.

Beauchamp, Tom L. and others, *The Human Use of Animals: Case Studies in Ethical Choice* (Oxford: Oxford University Press, 1998).

Beauchamp, Tom L. and David DeGrazia, *Principles of Animal Research Ethics* (Oxford: Oxford University Press, 2020), https://doi.org/10.1093/med/9780190939120.001.0001.

Bout, Henriëtte J., J. Martje Fentener van Vlissingen and Edgar D. Karssing, 'Evaluating the Ethical Acceptability of Animal Research', *Lab Animal*, **43**.11 (2014), 411–414, https://doi.org/10.1038/laban.572.

Brønstad, Aurora and others, 'Current Concepts of Harm-Benefit Analysis of Animal Experiments—Report From the Aalas-Felasa Working Group on Harm-Benefit Analysis – Part 1', *Laboratory Animals*, **50**.1S (2016), 1–20, https://doi.org/10.1177/0023677216642398.

Cherfas, Jeremy, 'The Scientist as Good Shepherd', *New Scientist*, **85**.1200 (1980), 1002–1003.

Cohen, Carl, 'The Case for the Use of Animals in Biomedical Research', *New England Journal of Medicine*, **315** (1986), 865–870. https://doi.org/10.1056/NEJM198610023151405.

Cohen, Carol, 'In Defense of the Use of Animals', in *The Animal Rights Debate*, ed. by Tom Regan (Lanham, MD: Rowman & Littlefield, 2001), 29–46. https://doi.org/10.5040/9798216418658.

Davies, Gail, 'Harm-Benefit Analysis: Opportunities for Enhancing Ethical Review in Animal Research', *Lab Animal*, **47** (2018), 57–58. https://doi.org/10.1038/s41684-018-0002-2.

Davies, Gail and others, *Review of Harm-Benefit Analysis in the Use of Animals in Research* (Animals in Science Committee, 2017). https://ore.exeter.ac.uk/repository/bitstream/handle/10871/31153/Review%20of%20harm%20benefit%20analysis%20in%20use%20of%20animals.pdf.

Donnelley, Strachan and Kathleen Nolan, 'Science, and Ethics', *Hastings Center Report*, **20**.3 (1990), 1–32, https://doi.org/10.2307/3563167.

Driver, Julia, 'The History of Utilitarianism', in *The Stanford Encyclopedia of Philosophy (Winter 2022 Edition)*, ed. by Edward N. Zalta and Uri Nodelman (Metaphysics Research Lab, Stanford University, 2022). https://plato.stanford.edu/archives/win2022/entries/utilitarianism-history/.

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

Gerritsen, Vanessa, *Güterabwägung im Tierversuchsbewilligungsverfahren* (Zurich: Schulthess, 2022).

Graham, Melanie L., and Mark J. Prescott, 'the Multifactorial Role of the 3RS in Shifting the Harm-Benefit Analysis in Animal Models of Disease', *European Journal of Pharmacology*, **759** (2015), 19–29. https://doi.org/10.1016/j.ejphar.2015.03.040.

Griffin, G. and others, 'Scientific Uses of Animals: Harm-Benefit Analysis and Complementary Approaches to Implementing the Three Rs', *Revue scientifique et technique (International Office of Epizootics)*, **33**.1 (2014), 265–272. https://doi.org/10.20506/rst.33.1.2283.

Grimm, Herwig, 'Turning Apples into Oranges? The Harm-Benefit Analysis and How to Take Ethical Considerations into Account', *ATLA*, **43**.2 (2015), 22–24. https://doi.org/10.1177/026119291504300211.

Grimm, Herwig and others, 'The Road to Hell Is Paved with Good Intentions: Why Harm-Benefit Analysis and Its Emphasis on Practical Benefit Jeopardizes the Credibility of Research', *Animals*, **7**.70 (2017). https://doi.org/10.3390/ani7090070.

Gutfreund, Yoram, 'Harm-Benefit Analysis May Not Be the Best Approach to Ensure Minimal Harms and Maximal Benefits of Animal Research – Alternatives Should Be Explored', *Animals*, **10**.2 (2020). https://doi.org/10.3390/ani10020291.

Hume, Charles W., 'The Ethics of Experiments on Animals', *Nature*, **167**.4241 (1951), 213–215. https://doi.org/10.1038/167213a0.

Hume, Charles W., *Man and Beast* (London: Universities Federation for Animal Welfare, 1962).

Kamm, Frances M., *Morality, Mortality Vol 1: Death and Whom to Save from It* (Oxford: Oxford University Press, 1993).

Kavka, Gregory S., 'The Numbers Should Count', *Philosophical Studies*, **36**.3 (1979), 285–294. https://doi.org/10.1007/BF00372632.

Kitcher, Philip, *Science, Truth, and Democracy* (Oxford: Oxford University Press, 2001). https://doi.org/10.1093/0195145836.001.0001.

Knight, Andrew, *The Costs and Benefits of Animal Experiments* (New York: Palgrave Macmillan, 2011), https://doi.org/10.1057/9780230306417.

Laber, Kathy and others, 'Recommendations for Addressing Harm-Benefit Analysis and Implementation in Ethical Evaluation – Report from the Aalas-Felasa Working Group on Harm-Benefit Analysis – Part 2', *Laboratory Animals*, **50**.1S (2016), 21–42, https://doi.org/10.1177/0023677216642397.

Lauwereyns, Jan, *Rethinking the Three R's in Animal Research* (Cham: Springer International Publishing, 2018). https://doi.org/10.1007/978-3-319-89300-6.

Lewis, C. S., *Vivisection* (Boston, MA: New England Anti-Vivisection Society, 1947).

Lyons, Dan, *The Politics of Animal Experimentation* (New York: Palgrave Macmillan, 2013).

Mason, Elinor, 'Value Pluralism', in *Stanford Encyclopedia of Philosophy (Summer 2023)*, ed. by Edward N. Zalta and Uri Nodelman (Metaphysics Research Lab, Stanford University, 2023). https://plato.stanford.edu/archives/sum2023/entries/value-pluralism/.

Mertz, Marcel and others, 'Interdisciplinary Animal Research Ethics – Challenges, Opportunities, and Perspectives', *Animals*, **14**.2896 (2024), 1–28. https://doi.org/10.3390/ani14192896.

Mill, John Stuart *On Liberty, Utilitarianism and Other Essays*, ed. By Mark Philp and Frederick Rosen (Oxford: Oxford University Press, 2015).

Moore, G. E., *Principia Ethica* (Cambridge: Cambridge University Press, 1922).

Müller, Nico D., 'Phase-Out Planning for Animal Experimentation: A Definition, an Argument, and Seven Action Points', *ALTEX*, **41**.2 (2024), 260–272. https://doi.org/10.14573/altex.2312041.

Niemi, Steven M., 'Harm-Benefit Analyses Can Be Harmful', *ILAR Journal*, **60**.3 (2019), 341–346. https://doi.org/10.1093/ilar/ilaa016.

Norcross, Alastair, 'The Scalar Approach to Utilitarianism', in *The Blackwell Guide to Mill's Utilitarianism*, ed. by Henry A. West (Blackwell, 2006), pp. 217–232. https://doi.org/10.1002/9780470776483.ch15.

Olsson, I. Anna S., Orsolya Varga, and Peter Sandøe, 'A Matter of Importance: Considering Benefit in Animal Ethics Review', *ALTEX Proceedings*, **4**.1 (2015), 33–36.

Orlans, F. Barbara, *In the Name of Science: Issues in Responsible Animal Experimentation* (Oxford University Press, 1993). https://doi.org/10.1093/oso/9780195119077.001.0001.

Phillips, Mary T. and Jeri A. Sechzer, *Animal Research and Ethical Conflict: An Analysis of the Scientific Literature 1966-1986* (New York: Springer, 1989). https://doi.org/10.1007/978-1-4612-3620-7.

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press

Pollo, Simone and Augusto Vitale, 'Evaluating the Scientific Uses of Animals: A Virtue-Consequentialist Approach for Harm/Benefit Analyses', *Journal of Applied Animal Ethics Research*, 2.2 (2020), 193–215. https://doi.org/10.1163/25889567-BJA10004.

Rawls, John, *A Theory of Justice* (Cambridge, MA: Belknap Press, 1999).

Regan, Tom, *The Case for Animal Rights* (Berkeley, CA: University of California Press, 2004).

Reiss, Julian and Biomedical Research. 'Philip Kitcher, Neglected Diseases, and Well-Ordered Science', THEORIA: Revista de Teoría, Historia y Fundamentos de La Ciencia, 24.3 (2009), 263–282. https://doi.org/10.1387/theoria.696.

Russell, William M. S. and Rex L. Burch, *The Principles of Humane Experimental Technique* (London: Methuen, 1959).

Sanders, John T., 'Why the Numbers Should Sometimes Count', *Philosophy & Public Affairs*, **17**.1 (1988), 3–14.

Sapontzis, S. F., 'On Justifying the Exploitation of Animals in Research', *The Journal of Medicine and Philosophy*, **13** (1988), 177–196. https://doi.org/10.1093/jmp/13.2.177.

Scanlon, T. M., *What We Owe to Each Other* (Cambridge, MA: Belknap Press of Harvard University Press, 1998).

Scarre, Geoffrey, *Utilitarianism* (London: Routledge, 1996).

Shaw, Jamie, 'On the Very Idea of Pursuitworthiness', *Studies in History and Philosophy of Science*, **91** (2022), 103–112. https://doi.org/10.1016/j.shpsa.2021.11.016.

Singer, Peter, *Animal Liberation: A New Ethics for Our Treatment of Animals* (Random House, 1975).

Slote, Michael and Philip Pettit, 'Satisficing Consequentialism', *Proceedings of the Aristotelian Society, Supplementary Volumes*, **58** (1984), 139–176. https://doi.org/10.1093/aristoteliansupp/58.1.139.

Sumner, L. W., 'Animal Welfare and Animal Rights', *Journal of Medicine and Philosophy*, **13**.2 (1988), 159–175. https://doi.org/10.1093/jmp/13.2.159.

Taurek, John M., 'Should the Numbers Count?', *Philosophy & Public Affairs*, **6**.4 (1977), 293–316.

Taylor, Katy D., 'Harms Versus Benefits: A Practical Critique of Utilitarian Calculations', in *The Ethical Case Against Animal Experiments*, ed. by Andrew Linzey and Clair Linzey (University of Illinois Press, 2018), 148–159. https://doi.org/10.5406/j.ctt2050vt5.20.

Timmermann, Jens, 'The Individualist Lottery: How People Count, But Not Their Numbers', *Analysis*, **64**.2 (2004), 106–112. https://doi.org/10.1093/analys/64.2.106.

Vogel, Isabel, *Espa Guide to Working with Theory of Change for Research Projects* (Ecosystem Services for Poverty Alleviation, 2012). https://www.espa.ac.uk/files/espa/ESPA-Theory-of-Change-Manual-FINAL.pdf.

Walker, Rebecca L., 'Beyond Primates: Research Protections and Animal Moral Value', Hastings Center Report, **46**.4 (2016), 28–30. https://doi.org/10.1002/hast.602.

Würbel, Hanno, 'Ethology Applied to Animal Ethics', *Applied Animal Behaviour Science*, **118**.3–4 (2009), 118–127. https://doi.org/10.1016/j.applanim.2009.02.019.

Cite this article: Müller, Nico Dario and Tristan Katz. 2026. Harm-Benefit Analysis for Animal Experiments Is Not Utilitarian. *Utilitas* **38**, 143–154. https://doi.org/10.1017/S0953820826100272

https://doi.org/10.1017/S0953820826100272 Published online by Cambridge University Press