Lives in the Balance:
Utilitarianism and Animal Research

ROBERT BASS
University of North Carolina at Pembroke

[This may differ in detail from the final published version in *The Ethics of Animal Research: Exploring the Controversy*, MIT Press 2012]

In the long history of moral theory, non-human animals – hereafter, just *animals* – have often been neglected entirely or have been relegated to some secondary status.[1] Since its emergence in the early nineteenth century, utilitarianism has made a difference by focusing upon happiness or well-being (and their contraries) rather than upon the beings who fare well or suffer. Inevitably, that has meant that human relations to and use of other animals have appeared in a different light. Some cases have seemed easy: once admit that the interests of animals matter and there can be little hesitation in condemning their cruel treatment. Among the more difficult cases has been the bearing of utilitarianism upon the use of animals in various kinds of research where, though the animals might suffer, there were believed to be prospects of great human benefit and where no cruel or malicious motives need be involved. What I shall provide in the current paper is an extended discussion of the bearing of utilitarianism upon practices of animal research. Since such practices have attracted both utilitarian criticism and defense, this will require the examination of arguments on both sides, including consideration of the human benefits, the animal costs, and the ways which one can be weighed against the other.

I. UTILITARIANISM AND ANIMAL RESEARCH: THE HISTORICAL AND THEORETICAL BACKGROUND

The historical connection of utilitarianism to animal research is both complex and disputed. It is common to ascribe a role to the early utilitarians in establishing the world's first animal protection society, Britain's Royal Society for the Prevention of Cruelty to Animals (RSPCA), and in the first anti-cruelty legislation in 1824.[2] Although animal research in some form was an ancient practice – Galen dissected animals for clues to human physiology – and the use of animals for crude toxicological tests is no doubt of antique vintage, it was not until the latter half of the nineteenth century that animal experimentation attracted widespread popular opposition. Perhaps, this was a matter of increasing scale and hence noticeability, since "the number of animal experiments conducted in Britain increased from 250 in 1881 (the first year that records were kept) to 95,000 in 1910."[3] Perhaps it had to do with the emergence of an increasingly large and prosperous middle class with leisure time to devote to causes beyond the immediate earning of a living. And perhaps utilitarianism was a factor. Utilitarian advocacy may have been cause, consequence or both, but it is surely safe to say that the same impulses appear to have been manifest both in utilitarianism and in the nineteenth century development of the anti-vivisection movement.

Moreover, regardless of the historical relation, the conceptual and theoretical bearings of utilitarianism upon animal research are worthy of investigation, given the importance of utilitarianism as an approach to moral theory and the world-wide scale of animal research. The

investigation requires something further to be said about how both animal research and utilitarianism are to be understood.

First, when we speak of animal research, let us limit ourselves to harmful and non-therapeutic research on animal subjects, for it is such research that most acutely raises questions of justification. Thus, we will set aside any research which is therapeutic in motivation, aimed at treating some disease or infirmity of the animal itself, and also any research which, whether therapeutic or not, is non-harmful to the test subjects, such as, perhaps, observational studies of their social lives.

Our present area of concern is animal research that is both non-therapeutic and harmful to the animals themselves. There is much of this kind, with perhaps 22 million animals killed annually for research in the United States and up to 100 million worldwide.[4] Three large categories can be distinguished. There is toxicological and safety testing, aimed at determining what substances are likely to be harmful to human beings.[5] There is biomedical research, aimed at discovering and determining the efficacy of treatments, preventive interventions and therapeutic options. Finally, there is the partially overlapping category of pure research, pursued for the sake of hoped-for gains in biological understanding, but which may or may not lead to other human benefit.

As for utilitarianism, it is most readily introduced in its classical form, as it was worked out by Bentham and Mill. It can be summed up in a few core propositions – that only happiness, understood as a function of pleasure, is good, that only unhappiness or suffering, understood as a function of pain, is bad, that quantities of happiness and suffering can be compared on a single scale and across individuals, and that the measure of rightness is the contribution of an action to the greatest overall happiness. The central idea is that the worth of an act is not to be determined by its conformity to some rule of duty or sacred principle; rather, rules of duty and sacred principles are to be determined by the worth of actions undertaken in accordance with them. And the worth of an act is to be judged in terms of its consequences – its impact upon overall happiness.

The theory has undergone much transformation since the time of Bentham and Mill, but key features have remained constant. Happiness may have been replaced with preference-satisfaction or more complex measures of well-being,[6] but the following features have remained.

First, utilitarianism is *consequentialist*: only contributions to good or bad consequences ultimately matter. If anything else matters for the utilitarian, like a character trait, it will be because of its contribution to the production of better consequences. Thus, a utilitarian might favor the disposition that leads a mother to care especially about the well-being of her own child, even when she could help other children more, on the grounds that the disposition pays off in better consequences on average, even if not in every case.[7]

Second, consequences are evaluated as better or worse in terms of the well-being of all affected. This is a kind of *universalism*: all of those whose welfare is at stake are considered.

Third, utilitarianism is *aggregative* in that it is assumed that benefits and harms can somehow be summed across individuals, so that an act which produces your modest pleasure and my great pain can be assessed as better *pro tanto* than an alternative that produces my great pain with no off-setting pleasure.

Fourth, it is *egalitarian* in that equal benefits and harms count equally, to whomever they accrue. From the moral point of view, my benefit will count for no more (and no less) than your equivalent benefit. It is, of course, understandable that I may be more attached to my benefit than

to yours, but, other things being equal, I should choose a greater benefit to you over a lesser benefit to me.

Fifth, utilitarianism is typically understood, by both defenders and critics, as a *maximizing* theory, as one that holds that good is to be maximized. Thus, we should act to produce the *best* consequences, so an objection to a policy can be founded upon the fact that some better option is available, even if the policy itself has good consequences.[8]

Refinements will concern us in due course, but a further key point is more temperamental than doctrinal: utilitarians have been, since the earliest days, *reformers*. They have been prepared to think through moral issues, assess their bearings on the welfare of all affected, and to draw the indicated practical conclusions. They were prepared to look upon the customary or traditional with a critical eye and to demand more defense than that things have long been done that way. When satisfactory defense was not forthcoming, they were prepared to demand reform or abolition.

Given this backgound, it is not surprising that animal research should have attracted suspicious attention from utilitarians. Their universalism meant that they were disposed to consider the well-being of animals. Their willingness to aggregate goods meant that in principle the animal well-being at stake in a particular action or practice might outweigh any human benefits derived therefrom. And, most importantly, the early utilitarians agreed that the happiness they wished to promote and the suffering they wished to prevent were not solely human prerogatives. There was, in other words, animal well-being to take into account.

In Bentham's much-quoted line, the important question when considering our treatment of animals is "Can they *suffer*?"[9] In a similar vein, Mill argued that "[t]he reasons for legal intervention in favour of children, apply not less strongly to the case of those unfortunate slaves and victims of the most brutal part of mankind, the lower animals," and that intervention should be based on "the intrinsic merits of the case," rather than upon "incidental consequences . . . to the interests of human beings."[10] Elsewhere, Mill is clear that the idea that animal pleasures and pains do not count compared to their human counterparts is on a par with such "superstitions of selfishness" as that the feudal nobility or slave-masters need not consider the interestsof serfs or of slaves.[11] And, in envisioning the utilitarian ideal as "an existence exempt as far as possible from pain and as rich as possible in enjoyments," he was explicit that such an existence should be "secured . . . so far as the nature of things admits, to the whole sentient creation."[12] Finally, Henry Sidgwick, the great nineteenth-century systematizer of utilitarianism, agreed that "it seems arbitrary and unreasonable to exclude from the [utilitarian] end . . . any pleasure of any sentient being."[13]

Contemporary utilitarians, notably Peter Singer, have been, if anything, even more explicit:

> If a being suffers there can be no moral justification for refusing to take that suffering into consideration. No matter what the nature of the being, the principle of [equal consideration of interests] requires that its suffering be counted equally with the like suffering . . . of any other being.[14]

Nor did Singer hesitate to apply this to animal experimentation, which was one of his two principal examples of speciesism – that is, of unjustifiable preference for human beings over other animals:[15]

> We tolerate cruelties inflicted on members of other species that would outrage us if performed on members of our own species. Speciesism allows researchers to regard the animals they experiment on as items of equipment, laboratory tools rather than living, suffering creatures.[16]

On the basis of the enormous and involuntary sacrifices imposed upon research animals, in comparison with the relatively few cases in which some outweighing human benefit can be found, Singer opposes practically all animal research.[17]

II. THE 'BENEFITS' ARGUMENT FOR ANIMAL RESEARCH

Interestingly, the principal argument in favor of animal research is also essentially utilitarian. Appeal is made to enormous benefits to human lives and well-being that are derived from animal research. For example, in "Animal Research is Vital to Medicine," Jack Botting and Adrian Morrison assert:

> Experiments using animals have played a crucial role in the development of modern medical treatments, and they will continue to be necessary as researchers seek to alleviate existing ailments and respond to the emergence of new disease.[18]

Throughout, their argument is that the medical benefits of animal research are substantial and that no satisfactory replacements are readily available. Thus, without the animal experimentation, those benefits will have to be foregone.

A problem for their case derives from what is missing. Like many defenders of animal research, they emphasize the benefits, but the lives and sufferings of research animals are conspicuous by their absence. For all that Botting and Morrison tell us, the research animals might as well be, as Singer suggested, just pieces of lab equipment. One cannot make a serious case that the benefits outweigh the harms, without counting the harms.

In similar terms, but with greater sophistication, the philosopher Carl Cohen argues that the medical benefits *do* outweigh the harms.

> [A] cogent utilitarian calculation requires that we weigh all the consequences of the use, and of the nonuse, of animals in laboratory research. . . . [A]n argument that is explicitly framed in terms of interest and benefit for all over the long run must attend also to the disadvantageous consequences of not using animals in research, and to all the achievements attained and attainable only through their use.
>
> The sum of the benefits of their use is utterly beyond quantification. The elimination of horrible disease, the increase of longevity, the avoidance of great pain, the saving of lives, and the improvement of the quality of lives (for humans and for animals) achieved through research using animals is so incalculably great that the argument of these critics, systematically pursued, establishes not their conclusion but its reverse: to refrain from using animals in biomedical research is, on utilitarian grounds, morally wrong.[19]

Certainly, this is the right *sort* of argument for the utilitarian proponent of animal research to be

making.[20] One has to be a bit suspicious, however, of benefits "incalculably great" and "utterly beyond quantification." In particular, if we are to carry out a utilitarian comparison with any hope of a determinate outcome, we will need at least to be assured, not only of the incalculability of the benefits, but that the harms suffered by research animals are not *also* incalculably great and unquantifiable.

What we appear to have are two opposed utilitarian arguments, both incomplete as they stand. The utilitarian critics of animal research point to the great harms undergone by animals, its utilitarian defenders to great human benefits. It should not be surprising that utilitarian arguments can be constructed on either side, at least so long as we have not further specified the relevant costs and benefits. In principle, utilitarians could go either way – favoring animal research if the weight of net benefit falls on that side, or opposing it, if the weight of net benefit falls differently.

To adjudicate, we need further empirical input, which I shall try to provide in the next four sections. In Section III, I will discuss some of what needs to be taken into account in a utilitarian assessment. In particular, what should figure in the assessment of *human* costs and benefits? In Section IV, I turn more directly to considering what it is reasonable to think the human benefits of animal research are. Then, in Sections V and VI, I consider what the animal harms are and how they are to be weighed in comparison with human benefit.

III. GETTING THE ACCOUNTING RIGHT

Surprisingly often, when the question of the human benefits of animal research is discussed, the obvious is overlooked. There are a variety of ways in which human benefits are overstated.

First, doing the accounting correctly means considering benefits *minus* costs. A list of benefits, however impressive, is not enough. Since there are two main properties animal tests are supposed to determine – therapeutic value and dangerous side-effects or toxicity, there are four kinds of mistakes to be concerned about, false positives and false negatives in each class. For therapies, we need to be concerned (1) about those that look promising with test subjects, but turn out not to be useful in clinical trials, and also (2) about those that would have been beneficial for us but for which no benefit was found in animal tests. And, with respect to dangerous effects, we need to be concerned both (3) about the cases in which a drug is dangerous to test subjects, though it would not have been to us, and (4) about the cases in which a treatment is safe for test subjects, but not for us.

Thus, any honest accounting of the human costs and benefits of animal research must estimate the human benefits *after* subtracting the losses in each of the above categories. Corresponding to the first, we have losses from ineffective treatments that looked promising on animals. Corresponding to the second and third, we have the continued suffering or premature deaths of those denied beneficial treatments that flunked an animal test. And corresponding to the fourth category, we have the losses of those whose harmful treatments passed the animal tests. The losses from beneficial treatments that flunked an animal test are especially difficult to estimate, since such treatments cannot legally proceed to clinical trials. But surely, it would be excessively optimistic to suppose that there are no such cases. The real benefits to human beings, once the losses are deducted, are bound to be a lot less than appears in, say, Botting and Morrison's catalogue of beneficial research. Since the detailed accounting is rarely done, it is difficult to tell how great the net benefits are, or even whether there are any.

Second, in order to fairly assess the benefits of animal research, we need to sort out the

effects traceable to that research from those we would have had anyway, sooner or later. Suppose we learn facts *A*, *B*, and *C* from animal research *X*. When we talk about the benefits of *X*, it is natural to count all the benefits of knowing *A*, *B*, and *C*. But that may be a mistake if, had we done non-animal research *Y* instead, we would, a year later, have discovered *B* and *C* anyhow. Then, the benefit that can *properly* be traced to animal research *X* is only the benefit of knowing *A* plus the benefit of knowing *B* and *C* a year earlier. It is not reasonable to count *all* the benefits of knowing *A*, *B*, and *C*, among the benefits of *X* unless we would be *deprived* of all of them if animal research *X* had never occurred.

Relatedly, we should include among the costs of animal research program *X*, those benefits of non-animal research program *Y* that will never see the light of day, since *X* is pursued in its place. Thus, if we would have learned facts *B*, *C* and *D* from research program *Y* then the benefits that would have flowed from knowing *D* should be counted as a cost of pursuing *X* rather than *Y*.[21]

These are not simply abstract possibilities. Non-animal research modalities exist, including methods for testing toxicity, simulating metabolic pathways, clinical and epidemiological studies, micro-dosing, *in vitro* testing, and computer and mathematical modeling.[22] It may be, as proponents of animal testing often maintain, that there are no fully adequate substitutes for animal research modalities. That, even if true, does not imply that without animal research no biomedical research would proceed or add usefully to our knowledge, nor does it imply that there are no benefits of non-animal research of which we *deprive* ourselves by pursuing animal research modalities instead. In estimating the benefits of animal research, we must deduct both those we would have anyway without the animal research and also the *value* of those we would have instead.

Third, animal research competes with other causes for societal support. At any given time, there is a limited total supply of resources that can be devoted either to medical research or to other worthy causes. Obviously, the utilitarian will favor using those resources in whatever ways have the maximum impact on saving lives, preventing disease and promoting the well-being of all affected. That may well mean that those resources should *not* be devoted to animal research, *even if we know it promises human benefit*.

For there is still a moral question to be raised: The animal research program will cost something, a certain number of dollars, say, to save a certain number of lives. Before concluding that we've gotten a good deal, we need to ask if we couldn't have achieved as great or greater gains for the same money, spent some other way. There is not an *a priori* answer, but in a world in which we know that much serious illness is due to dietary and lifestyle choices, it seems that nutritional or lifestyle education might well produce greater returns than animal research.[23] So might many other options, since much risk to life and health is due not to our ignorance about what to do but to our failure to take steps that we have every reason to believe would be effective.

In fact, there is a good utilitarian case that not much in the way of scarce resources should go to support medical research *at all*, or for that matter to medical *treatment* for citizens of the wealthier countries of the world. The reason is that most research and treatment goes to providing therapies and interventions that marginally extend life or enhance quality of life for people who are already, by historical and current world-wide standards, exceptionally fortunate. For the foreseeable future, there are and will continue to be millions of children in the less developed countries of the world who can be given a chance at decent lives at far less expense

than adding another six months or two years to the end of already long lives. Whether in this way or some other, when greater returns are available than from animal research, that will have to tip the utilitarian scale.

IV. THE EFFICACY OF ANIMAL RESEARCH

When the question is fairly considered and errors are avoided, how does animal research perform as an engine of human benefit? Much is controversial here, and there are specialists far better qualified than I to address the issue,[24] some of them contributors to the present volume. For the most part, I shall leave the presentation of the case to them. However, I can hardly proceed without some information about the value of animal research. I think that can be provided while avoiding most technicalities, since animal research itself has recently come under scientific scrutiny.

Consider the fact that most institutions in the United States that engage in animal research screen the research through internal Institutional Animal Care and Use Committees (IACUCs).[25] The job of an IACUC is to evaluate proposed animal research both for its scientific merit and for its ethical acceptability. On the scientific front, they are supposed to provide oversight to help ensure that scientifically important work is being done, which includes the requirement that it not simply duplicate work done elsewhere.[26] On the ethical front, the IACUC's mandate is to make sure that the procedures to be followed will minimize animal harm, pain and distress so far as is consistent with sound scientific practice, and that there are no scientifically acceptable alternatives to the animal protocols – that the research cannot be pursued using fewer animals, causing less pain and suffering, or using non-animal research modalities. All of these are legal responsibilities of the IACUC.[27]

IACUC members are expected to have relevant expertise, as well as to fulfill a representative function. By USDA standards, IACUCs must have no less than three members, including one veterinarian and someone unaffiliated with the institution. By the more restrictive Public Health Service standards, there must also be at least "[o]ne practicing scientist experienced in research involving animals," someone with a non-science background, such as law or ethics, and someone to represent general community interests.[28]

Underlying assumptions in justifying the role of IACUCs in providing scientific and ethical oversight are that there are determinable answers to the ethical and scientific questions, that the expert members will generally be deploying the same standards, and thus, that they should be expected to reach substantially concordant conclusions. The point can be illustrated with any kind of expertise. Experts on chess will largely agree on what is a good move. Experts in mathematics will largely agree on what is or is not a rigorous proof. And so on.

If we apply this to IACUCs, we should expect significant agreement both on what research is ethically acceptable and on what is scientifically promising. However, a study published in *Science*, comparing the judgments of different IACUCs found that when they were presented with identical research protocols, stripped of identifying information, "protocol reviews did not exceed chance levels of intercommittee agreement."[29] In 79% of the research protocols reviewed, the second committee reached a different conclusion than the first.

Since different committees agreed no more than would be expected by chance, there are serious questions about the expertise being invoked. If we assume substantial agreement on the scientific merits of animal protocols, chance levels of overall agreement imply deep divisions in ethical assessment. Since the ethical standards are largely dictated by applicable law and

regulation, it may be more reasonable to assume substantial agreement on the ethical acceptability of various animal research protocols.[30] But then, chance levels of agreement imply either that there is no such thing as scientific expertise in evaluating the merits of research protocols or that there is no tendency for the composition of IACUCs to reflect such expertise. Whatever the explanation, whether it be in terms of ethical or scientific disagreement, or some combination, the result is troubling. Since an institutional safeguard, designed to insure the ethical acceptability and scientific merit of animal research, appears to be performing no better than chance, there is little reason to suppose that research protocols approved by IACUCs are either scientifically or ethically acceptable.

A possible reply would be that the evidence is consistent with the hypothesis that most animal protocols submitted to IACUCs are of very high quality, and thus that the differences between them are sufficiently subtle that it is unsurprising if experts disagree. Experts disagree on questions at the frontiers of physics and other sciences, too, not because there is no relevant expertise, but because the evidential considerations are subtle, less than fully conclusive and are difficult to articulate. Similarly, that *might* be the explanation for low levels of concordance between IACUCs.[31]

On the face of it, that does not appear likely. In particular, it fails to specify any mechanism to insure the high quality of research proposals. Without some such mechanism we should expect a wide range in quality. Fortunately, we do not have to confine ourselves to guesswork. If the hypothesis that the research proposals considered are all of high quality is correct, then we can predict that IACUC-approved research will turn out, by independent criteria of assessment, to be of uniformly high quality. As we shall see, this is not the case.

What quality and what human benefit do we find in animal studies? Recently, Pablo Perel and his colleagues took a step toward answering that question. They published, in *BMJ*, the *British Medical Journal*, a systematic review of animal studies relevant to treatment effects known from clinical trials. Six clinical interventions were selected, and then, an extensive search was conducted for all relevant animal studies. In three of the cases, the animal studies corresponded to the clinical trials, and in three they did not. Moreover, the authors repeatedly found that the "animal studies were of poor quality." In the case of one intervention, they also found strong evidence of publication bias.[32] Though the authors emphasize that an examination of six interventions is not enough to pass judgment upon all animal research, they conclude that the "[d]iscordance between animal and human studies may be due to bias or to the failure of animal models to mimic clinical disease adequately."[33]

Proceeding from the other end, Daniel Hackam and Donald Redelmeier published in *JAMA*, the *Journal of the American Medical Association*, details of a study of the translation of animal research into clinical results. They examined articles from seven leading scientific journals reporting animal research, ranked by number of citations – which is to say, by their post-publication influence – to select 2000 highly cited animal studies. Out of these, they sought animal studies that examined preventive or therapeutic interventions and followed up with a literature search to find clinical or human trials of the same interventions. They sought to examine "how often highly cited animal studies translate into successful human research."[34] Their finding was: *not very often*. Even though the animal studies included could reasonably be said to be the best of the best, with a median of 889 citations each, only 37% of the results were replicated in human trials. In summing up, the authors note that the low translation rate may be unrepresentatively *high*, since they only examined highly cited studies from leading journals,

which is to say, studies which had already undergone rigorous peer review and had attracted significant scientific esteem. It is very plausible that less rigorously reviewed or less cited studies would also be less likely to translate well to clinical applications. Hackam and Redelmeier close with the caution that "poor replication of even high-quality animal studies should be expected by those who conduct clinical research."[35]

It would surely be going too far to argue on the basis of these studies that there is no human benefit to be derived from animal research, but they suggest substantial reasons for caution. Institutional Animal Care and Use Committees do not exhibit more intercommittee agreement than would be expected by chance, so there is little reason to think that they are approving only, or even primarily, well-designed animal studies. This is confirmed by the Perel study, which found poor methodological quality and lack of concordance between animal studies and clinical trials. Further confirmation comes from Hackam's and Redelmeier's much more extensive review of the correspondence between animal studies and clinical trials: correspondence was poor, even for the best-designed and most influential animal studies.

One has to ask, with the authors of another recent study, "Where is the evidence that animal research benefits humans?" As Pandora Pound and her associates introduce the issue:

> Clinicians and the public often consider it axiomatic that animal research has contributed to the treatment of human disease, yet little evidence is available to support this view. Few methods exist for evaluating the clinical relevance or importance of basic animal research, and so its clinical (as distinct from scientific) contribution remains uncertain. Anecdotal evidence or unsupported claims are often used as justification—for example, statements that the need for animal research is "self evident" or that "animal experimentation is a valuable research method which has proved itself over time." Such statements are an inadequate form of evidence for such a controversial area of research. We argue that systematic reviews of existing and future research are needed.[36]

Plainly, if systematic reviews are needed to overcome the weakness of merely anecdotal evidence and unsupported claims, the matter is not already settled. Certainly, it is not settled by proclamations that the serious debate is already over.

What, then, is a reasonable conclusion? Certainly, there are many who think the human benefits of animal research are large. Equally, there are some who think it clear that there are no benefits, or even that animal research is positively harmful in its impact upon human health and well-being. For my purposes, I see no need to make either of the latter claims. No part of my argument depends upon there being either no benefits or net harm.

But if the human benefits were large and clear, it would be difficult to explain why they have proven so hard to find in careful reviews of the correspondence or lack thereof between animal studies and clinical trials or why methodology in animal research is so often poor.[37] Both classes of facts appear explained, at least in part, if the bodies responsible for approving animal research, the IACUCs, do not reliably distinguish between what is scientifically valuable or dubious. The most reasonable working hypothesis is that the human benefits of animal research are either small or unclear.

V. WEIGHING BENEFITS AND HARMS

How do we assess the relative weights of benefit and harm from animal research, counting the welfare of both humans and other animals? Two preliminary issues need to be addressed: how actual harms are to be weighed against possible benefits and whether there is any relevant harm to animals.

The first issue concerns the frequent rhetorical attempts to pose the issues of animal research as simple oppositions between, say, a child and an animal: Would you sacrifice a dog to save a child? Posed in this way, most people would favor the child. So, arguably, would most utilitarians. They would appeal to the child's greater life expectancy and the greater richness of experience and enjoyment open to the child as she grows into adulthood. But that is not the actual choice that faces us.

In reality, on one side of the ledger, we have *possible* benefits for human beings – lives that might be saved or prolonged, diseases that might be prevented, cured or ameliorated. On the other side, we have *actual* harms. Of the laboratory animals, many will suffer, and almost all will be killed. We do not face the relatively easy choice between the child and the dog. In considering animal research, we face the much harder choice between large numbers of animals, who will surely suffer and die, and unknown numbers of human beings who may or may not benefit. The death and suffering is assured; the benefit is not.

How should actual harms be compared to merely possible benefits? For the moment, let us focus on the harms and benefits, while leaving aside matters specific to comparing animals with human beings. We do not want to simply treat similar harms and benefits as equal in importance, when some of them are actual and others are merely possible. Giving up 'eyes for possible eyes, teeth for possible teeth, lives for possible lives' is less than appealing, since it is apt to reduce the actual supply of eyes, teeth and lives. There must be some discounting, whereby possible benefits count for less than actual benefits, and actual harms for more than possible harms.

In principle, it is uncontroversial how to do this.[38] We compare the *expected values* of harms and benefits, where their expected values are equal to their actual values discounted by their probabilities. Thus, the expected value of an actual harm will be the same as its actual value, since it will be the actual value multiplied by a probability of one hundred percent. Expected values of uncertain future benefits will be less than the values of corresponding actual benefits, since they will be discounted by their probability. A fifty percent chance of a benefit is worth only half as much as the actual benefit. And even when expected values are equal, for example, with a chance to sacrifice a life for a fifty percent chance of saving two, the question could be pressed: why sacrifice a life to gain something worth no more than a life? There would be no clear justification for ending one life, without the prospect of something *greater* in value, such as a fifty percent chance of saving three or five or ten lives.[39]

There is an immediate application to animal research. Given that experiments certainly cause animal death and suffering but do not certainly produce human benefit, the outweighing expected value can only be had if there are *considerably* greater human benefits in prospect than animal harms. And for there to be considerably greater benefits, there must be, more simply, great benefits.

The calculation is further complicated by the fact that, to this point, I have idealized by assuming we have accurate probability information. When accurate probability information is available, we know by how much we have to discount future possible benefits in order to

compare them with actual harms. But in fact, nothing of the sort is available. On a case by case basis, there is virtually no reasonable expectation of benefit to human health or well-being, much less outweighing benefit.[40] If there is to be outweighing benefit, it must be the benefit derived from, at least, broad swaths of animal research over long periods of time. For reasons indicated earlier, we have little evidence of this sort. It may be that animal research yields substantial human benefit, but too many relevant questions, both about the benefits and about what would happen instead, go mostly unasked or unanswered. We know neither how large the benefits are nor what their probability is. Since the best working hypothesis is that the human benefits of animal research are either small or unclear, we are not in a position to claim justification. If the benefits are small, they cannot outweigh large harms to animals; if they are unclear, even as to probabilities, then we do not know there to be any outweighing benefit.[41]

The second preliminary issue arises from the fact that there would be no need for large, outweighing benefits if there were no harms to outweigh. What makes animal research look morally problematic is that it causes suffering and death to animals for the sake of benefits they don't understand, have no reason to care about, and to which they are not able to consent. That seems hard to justify. *Prima facie*, anyone trying to justify it has the burden of proof.

The burden could be sidestepped, however, if it were supposed that the animals have no conscious awareness, no feeling or capacity to suffer. Some defenders of the *status quo* with respect to animal research evidently find the thought tempting. Notable is Peter Carruthers, who apparently thinks (somewhat tentatively) that "the mental states of animals are non-conscious," and that animals "cannot be appropriate objects of moral concern."[42] This is the kind of position that could be adopted only in a desperate attempt to save a theory. It is true that *if* lab animals were insentient beings, incapable of experiencing the world or caring what happened to them, then it would be hard to see what moral importance their lives, in or out of the laboratories, would have. But that view of animals flies in the face, not only of common sense, but of the combined weight of evidence from several sciences, including ethology, neurophysiology, and evolutionary biology. The denial of animal consciousness is not an option that can be taken seriously by people who lay claim to scientific respectability.[43]

This may seem sufficient to bring in a negative utilitarian verdict on animal research. There are large harms to the animals that do not appear to be matched by large prospective benefits for humans. Questions remain, however, having to do with the comparison of animal losses to human gains. How are we to assess the relative worth of human and animal suffering, or of human and animal enjoyment? Differences in human and animal minds might matter here. If a human being is more exquisitely aware, perhaps what a human being undergoes in a single minute of intense suffering is greater than what another animal undergoes in many hours. Perhaps, even modest human enjoyment is richer and deeper than anything another animal could ever experience. If such were the case, there might be no need for a large benefit – or for what appears to us as a large benefit – to outweigh significant animal harm. A little human benefit might go a long way.

This is difficult territory, akin to debates about interpersonal utility comparisons among humans,[44] but we can make progress. We should acknowledge immediately that there is much that we do not know and that our insight into animal minds is partial at best. That does not mean we are completely at a loss.

Before turning to our knowledge of animal experience, let us begin with what is problematic about utility comparisons among human beings. Roughly, utility is a measure of the

satisfaction of preferences, and, arguably, we have no better way to characterize the well-being at which the utilitarian aims than in terms of utility.[45] The problem is that the measure of Janet's utility is anchored to her preferences and of David's utility to his, so how are we to say which is the utility-maximizing option if we must choose between, say, Janet suffering a pinprick and David having his hand mangled?

There are real questions about how to formally represent the problem, but part of the point of the example was to suggest that, whatever the difficulties of formalizing our judgment, we are not really in the least doubtful that it is worse that David has his hand mangled than that Janet suffers a pinprick. Nor are we puzzled whether charitable contributions should be dispatched to random millionaires on the ground that, for all we know, that might produce more utility than contributions directed to the poor or starving. Certainly, there is room for doubt about our access to other *human* minds, but it is limited and partial, rather than all-encompassing, doubt. The key point is to avoid skepticism about other minds.

Now, I do not have any ready solution to the problem of our knowledge of other minds to outline in a paragraph or two, but that is not necessary, so long as we remember that, however puzzling knowledge of other minds is in theory, we are not actually doubtful that we can reach reliable judgments in many cases. Other things being equal, we know that it is worse to have a hand mangled than to suffer a pinprick, that hundred dollar donations are better used to feed the starving than to line the pockets of the wealthy, and much more. This is the kind of knowledge of other minds, judgments of which we are confident in our daily lives and interactions, that can be leveraged into addressing how harms to animals are to be weighed against human harms and benefits.

VI. RETHINKING SINGER'S CHALLENGE

With this in mind, consider Singer's challenge, in response to the question, "[w]ould we be prepared to let thousands of humans die if they could be saved by a single experiment on a single animal?"[46] Singer's reply is to ask if the experimenter would be willing to use a human infant, perhaps one seriously handicapped, so that there is no prospect that it will ever have any level of mentality beyond that of a lab animal, with equal prospect of benefit. If the experimenter would not consent to use a human subject in appropriately parallel circumstances, he must be "willing to use a nonhuman animal only because he gives less consideration to the interests of other species than he gives to members of his own – and this bias is no more defensible than racism or any other form of arbitrary discrimination."[47]

This is relevant in the context of Singer's argument, but it can be put to different use. Utilitarianism does not privilege one species over another; what is valuable is happiness or well-being, no matter whose it is. To that extent, there is no need to make Singer's argument against speciesism. What we can consider, however, is the question whether we should conduct medical research upon severely handicapped humans, that is, upon humans with no greater intellectual or emotional capacity or potential than laboratory animals. Since they are members of our own species, the medical benefits would probably be greater than could be expected from animal research. Certainly, there is no reason to expect the benefits to be less. But most of us are convinced that such treatment would nevertheless be wrong. If we are utilitarians, that has to mean that we are convinced that the benefits to be derived from human experimentation do not outweigh the costs, primarily in the form of suffering and death imposed upon the test subjects.

Now, on the assumption that this judgment is reasonable, we have what we need to make

progress on weighing animal interests against human interests.[48] Even if normal human minds are much more richly aware than normal animal minds, the same can be said of the comparison between normal human minds and those of the seriously handicapped. We are confident enough of our knowledge of the mental lives of the handicapped that we are unwilling to experiment on them for the benefit of the rest of us. We should be equally confident when animals are the test subjects.

VII. RECASTING THE DEBATE: ACTS VS. RULES

Utilitarians are not absolutists in their concrete moral views. They will not hold that killing or lying or animal research are always wrong (or always right), no matter what.[49] Of course, they will maintain that it is always better to maximize utility, but that abstract prescription provides insufficient guidance in the absence of further information about what, concretely, does contribute to greater utility. The utilitarian attitude is that "we should always try to find out as much as possible about the probable consequences of our actions. Without this information, our decisions about what we ought to do should be subject to revision in the light of more complete information."[50]

This means that, however good the utilitarian case against animal research in general, it will be possible in principle to find cases in which it seems justified. When and if human benefit seems procurable at a sufficiently low price in animal harm, utilitarians will not object. Peter Singer was recently in the news for acknowledging this in a BBC documentary. In conversation with Oxford neurosurgeon, Tipu Aziz, he agreed that helping thousands of sufferers of Parkinson's disease by way of research with less than a hundred monkeys would be justifiable, at least if the facts were as Aziz represented them and if there were no other way to make the necessary discoveries. On one hand, this drew criticism from some animal advocates, who wanted nothing less than an absolute prohibition and apparently felt betrayed. On the other, it occasioned praise from defenders of the *status quo* for realism and honesty, who read it as an overdue admission of the value of animal research. In fact, it appears that neither side understood him well, because his position has not changed. As a utilitarian, he judges actions by their consequences, including consequences for animals, and is very doubtful about most animal research, but he does not rule out the possibility that some might have good enough consequences, on balance, to be justified.[51] For anyone in doubt where he stands, Singer later wrote:

> Even if some individual experiments may be justified, this does not mean that the institutional practice of experimenting on animals is justified. Given the suffering that this routinely inflicts on millions of animals, and that probably very few of the experiments will be of significant benefit to humans or to other animals, it is better to put our resources into other methods of doing research that do not involve harming animals.[52]

The episode brings into focus an important question. Utilitarians look for prospective justification. In order to decide what to do, we need information, the best we can get, about which options promise to pay off in better consequences. Perhaps, occasionally, animal research will credibly make that promise, even taking animal interests into account.

That is not an argument that we can safely institutionalize animal research. Given the expenses of supporting laboratories, facilities, researchers, support staff, and animals, we will not

be able to authorize all and only that portion of animal research that is prospectively justified on a case-by-case basis. Instead, if much animal research is going to be done, there will have to be large educational and research organizations, grant-issuing bodies, approval boards, and so forth. In fact, there will have to be an institutional matrix much like what we have now.

After-the-fact justifications of instances of animal research are not enough. In order to guide practice, we need to be able to identify valuable research in advance. Neither, however, is it enough if we are able, before the fact, to justify particular experiments or research programs, given that animal research will not be done piecemeal. The real question that faces the utilitarian is whether we can *institutionalize* animal research in a way that is prospectively justified. That is, can we authorize large quantities of animal research in a way that (a) has a high probability of supporting research that itself has a high probability of producing net benefits, and that (b) has a low probability of supporting research that does not have good prospects of producing net benefits?

On present evidence, the needed prospective justification seems unlikely. First, as we have seen, we do not have strong evidence for large human benefits of animal research in any case. At the very least, much better and more detailed accounting is needed to make the case for substantial human benefit. Second, since the evidence indicates that Institutional Animal Care and Use Committees do not reliably discriminate between promising and unpromising science, we would need to develop new and more rigorous evaluation procedures. It is difficult to know just how to proceed, however, given that our best efforts so far, by way of designing and regulating IACUC procedure, have proven no more reliable than flipping a coin. Third, we still need to consider the magnitude of animal harm. It will not be enough to find a small human benefit due to animal research. There is a great deal of harm to animals at stake: imposition of fear, pain, distress, injury, confinement, suffering, disease and death upon many of the test subjects. Only enormous and very clear prospective benefits to humans could outweigh that. In the absence of that kind of evidence, it is hard to see any utilitarian justification for institutionalized animal research.

NOTES

1. For a reminder of some of the ways in which this is less than the whole story, see Sorabji 1993.
2. Technically, the RSPCA only became “Royal” in 1840; prior to that it was only the SPCA. (Information from the RSPCA website at http://www.rspca.org.uk, accessed 22 January 2007.)
3. Perry, *et al.* 2005, 18.
4. Perry, *et al.* 2005, 7. It is difficult to obtain reliable estimates. In the United States, for example, there appears to have been no official count since 1986 that includes the birds and especially the rodents, who make up the majority of animal experimental subjects (85 to 95 percent). The 1986 count, conducted by the Congressional Office of Technology Assessment, concluded that 17 to 22 million animals were killed for research annually. See Stephens 2002.
5. Much toxicological and safety testing is not of medical procedures but of consumer products. There are few defenders of such animal testing, since there are neither legal requirements nor a lack of other effective ways to assure consumer safety.

6. See, for example, Hare 1981 and Brink 1989, Chapter 8.
7. Hare 1981, 135-140.
8. There are non-maximizing versions of utilitarianism, when it comes to the theory of the right or obligatory. In place of identifying rightness with production of the best, *satisficing* theories require the production of what is *good enough*, with that standing in need of further articulation and specification. See Slote 1985, especially Chapter 3. *Scalar* theories deny that there is some point or property that identifies rightness. Instead, actions can *only* be ranked as better or worse in terms of consequences, and, of course, the better an action, the more reason an agent has to take it. See Howard-Snyder and Norcross 1993, Norcross 1997 and Norcross 2006. For non-maximizing varieties of utilitarianism, the point from the text would need reformulation or might fail to apply. A scalar theory would allow criticism of an option for being less good than some alternative, but would not take that fact to show that the inferior option is impermissible. A satisficing theory might or might not allow criticism on the same basis, depending upon whether the option in question qualifies as good enough. More generally, since satisficing theories are indeterminate in their implications as long as a threshold for satisficing is unspecified and arguably morally arbitrary as well (Howard-Snyder and Norcross 1993, 111-113), I shall assume for the remainder of the chapter that we have to do with maximizing or scalar versions. Both will agree on how options and states of affairs are to be ranked *vis-à-vis* one another; they will agree that there is reason to select better options and bring about better states of affairs, while disagreeing only over whether the deontic predicates, such as *required* and *impermissible*, have any application.
9. Bentham 1823, 381.
10. Mill 1965, 952.
11. Mill 1969b, 186-187.
12. Mill 1969a, 40-41.
13. Sidgwick 1981, 414.
14. Singer 1975, 8.
15. Singer 1975, 23-24, 27-91.
16. Singer 1975, 62.
17. "If it really were possible to save many lives by an experiment that would take just one life, and there were *no other way* those lives could be saved, it might be right to do the experiment. But this would be an extremely rare case." Singer 1975, 78.
18. Botting and Morrison 1997, 83.
19. Cohen 1986.
20. Cohen does not appear to be a utilitarian himself, but he is endorsing the cogency of a utilitarian defense of animal research.
21. I am assuming that *D* would not have been learned without pursuing program *Y*.
22. For relevant material, consult the Center for Alternatives to Animal Testing at Johns Hopkins University <http://caat.jhsph.edu/index.htm> and the European Centre for the Validation of Alternative Methods <http://ecvam.jrc.it/index.htm>. See also the discussion of alternatives in Mukerjee 1997.
23. The American Cancer Society, for example, estimates that about a third of all cancers are connected to "poor nutrition, physical inactivity, overweight, and obesity" with many more due to smoking, and "at least half of all cancer deaths could in principle be avoided

by the application of existing knowledge." See <http://www.cancer.org/downloads/STT/CPED2006PWSecured.pdf>, accessed 29 January 2007.

24. I would especially recommend Greek and Greek 2002 and LaFollette and Shanks 1997.
25. Similar institutional arrangements are in effect in many other countries.
26. The non-duplication requirement could, of course, also be represented as a factor in assessing the ethical acceptability of the protocol.
27. See the *Institutional Animal Care and Use Committee Guidebook* (2002), especially 85-119, <http://grants.nih.gov/grants/OLAW/GuideBook.pdf>, accessed 15 January 2007.
28. *Institutional Animal Care and Use Committee Guidebook* (2002), 14, <http://grants.nih.gov/grants/OLAW/GuideBook.pdf>, accessed 15 January 2007.
29. Plous and Herzog 2001.
30. A complication is that judgments of ethical acceptability may vary, depending upon judgments of scientific merit. A protocol that causes a given quantity of animal suffering might be judged acceptable for research with great scientific promise but not when the scientific promise is less.
31. This line of reply works, to whatever extent it does, if animal protocols submitted for review fall into *any* narrow band of differential quality. Chance levels of agreement would be no more surprising if proposed protocols were uniformly of *low* quality (or uniformly of mediocre quality, etc.), provided that IACUCs are generally disposed to approve at least some proposals.
32. Statistical evidence of publication bias in connection with the other interventions could not be assessed due to the smaller numbers of studies included, though it was not ruled out. Perel, *et al*. 2006, 4.
33. Perel, *et al*. 2006, 1.
34. Hackam and Redelmeier 2006, 1731.
35. Hackam and Redelmeier 2006, 1732.
36. Pound, *et al.* 2004, 514. Endnotes omitted.
37. Perel, *et al*. 2006, 2. Is poor methodology the cause of poor correspondence? Perhaps, but Hackam and Redelmeier 2006 shows poor correspondence, even with much better methodological quality.
38. At least, it is uncontroversial on the assumption that it is morally appropriate to compare and trade off benefits and harms between different individuals. That assumption is itself controversial, but should not be controverted in the course of a utilitarian argument over the relative weights of animal harms and human benefits in animal research. For some technical details on expected value, see, e.g., Luce and Raiffa 1985, especially §2.4, and Heap, *et al.* 1992, 3-11.
39. There is also the complication that this only seems reasonable if it is part of a large set of gambles. If it is only going to happen once, the unfortunate victim is surely entitled to complain that he will certainly die, while there is only a fifty-percent chance of anyone being saved.
40. According to John Rennie, editor of *Scientific American*, "The short-term benefits of most experiments are virtually nil," but he adds that "the long-term benefits are incalculable. How do we enter them in a utilitarian ledger?" Perhaps, Rennie means that the benefits are incalculably *large*, but then that would not pose much problem for utilitarian calculation, at least if the harms could be confined to the calculably small.

There is evidence this is not his meaning. The quote I gave is embedded in a longer passage:

> In my opinion, the arguments for banning experiments on animals--that there are empirically and morally superior alternatives--are unpersuasive. And even some of the moral philosophies favoring reduced use of animals offer little in the way of real guidance. Utilitarians, for instance, ask that the suffering of animals be counterbalanced with good results. But that principle is unmanageably subjective and may even be prejudiced against research realities. The short-term benefits of most experiments are virtually nil, and the long-term benefits are incalculable. How do we enter them in a utilitarian ledger? Does increasing the sum of human knowledge count as a good? (Rennie 1997, 4)

The argument that considering the suffering of animals is "unmanageably subjective" raises serious questions. If it means that there is no respectable way to determine how much animals suffer in the course of experimentation or whether the benefits outweigh the suffering, then, surely, that's an argument against the animal experimentation. Saying that such consideration may be "prejudiced against research realities" suggests that the problem is not that consideration of animal suffering introduces unacceptable subjectivity, but that much research would not pass the test if animal suffering seriously had to be taken into account. Further indication that this is the real issue comes in the attempt to find something else, "increasing the sum of human knowledge," to add to the pro-research side of the ledger.

41. Perhaps, though, the assumption of large animal harms is not in order. This is further addressed below.
42. Carruthers 1992, 192. In fairness to Carruthers, this is not his only argument that animals lack the kind of moral significance that would limit their use as research subjects.
43. Evelyn Pluhar provides extensive and persuasive rebuttals both to Carruthers and to others who have denied animal consciousness. (Pluhar 1995, 12-54). Excellent additional discussion of Carruthers' position and arguments, together with insightful and scientifically informed discussion of the issues pertaining to animal pain, can be found in Allen 2004. That animals not only are our fellow-sufferers, but also share with us a rich variety of capacities for enjoyment and pleasure is persuasively argued in Balcombe 2006.
44. For very interesting discussion, see Nozick 1985 and Binmore 2005, 113-128.
45. For an attempt of this kind, see Railton 1986a and 1986b.
46. Singer 1975, 74.
47. Singer 1975, 76. See also surrounding text.
48. R.G. Frey is a utilitarian who denies the judgment:

> My position is that *if* we continue to employ an argument from benefit to justify animal experimentation in medicine, and *if* we continue to employ a discussion of the comparative value of human and animal life to justify using the animal in preference to the human, then, because not all human life has the same value as normal adult human life, we are faced with the prospect of human experiments, of, that is, using some humans as we currently use animals. The antecedents are widely accepted in the medical research community; with that the case, I see no way to avoid the conclusion. (Frey 1997, 131)

49. They will not hold this even if they are rule or indirect utilitarians of some stripe. First, they may include exceptions within rules justified on utilitarian grounds. The rule will not be, say, that killing is always wrong, but rather that killing is wrong unless some carefully circumscribed exception, such as self-defense, is relevant. Second, the rules themselves will be sensitive to available evidence. They will not hold no matter what. Should it turn out that some apparently justified rule has disutilitarian consequences, that will be a reason, other things being equal, for altering the rule.
50. Singer 1980, 328.
51. See Jaschik 2006 and Singer 2006, and also note 17 above, where Singer's position, as of 1975, is cited.
52. See <http://www.princeton.edu/~psinger/faq.html> (accessed 14 February 2007).

BIBLIOGRAPHY

Allen, Colin. "Animal Pain." *Noûs* 38 (2004): 617-643.

Anderson, Lynn C., *et al. Institutional Animal Care and Use Committee Guidebook (2nd ed.)*. 2002. <http://grants.nih.gov/grants/OLAW/GuideBook.pdf > (15 January 2007).

Balcombe, Jonathan. *Pleasurable Kingdom: Animals and the Nature of Feeling Good*. London, UK: Macmillan, 2006.

Bentham, Jeremy. *Principles of Morals and Legislation*. In *The Utilitarians*. Garden City, NY: Anchor Books, 1823.

Binmore, Ken. *Natural Justice*. Oxford, UK: Oxford University Press, 2005.

Botting, Jack H., and Adrian R. Morrison. "Animal Research is Vital to Medicine." *Scientific American* 276 (1997): 83-85.

Brink, David. *Moral Realism and The Foundations of Ethics*. Cambridge, MA: Cambridge University Press, 1989.

Carruthers, Peter. *The Animals Issue*. Cambridge, MA: Cambridge University Press, 1992.

Cohen, Carl. "The Case for the Use of Animals in Biomedical Research." *New England Journal of Medicine* 314 (1986): 865-869.

Frey, R. G. "Moral Community and Animal Research in Medicine." *Ethics and Behavior* 7 (1997): 123-136.

Greek, C. Ray, and Jean Swingle Greek. *Specious Science*. New York, NY: Continuum, 2002.

Hackam, Daniel G., and Donald A. Redelmeier. "Translation of Research Evidence from Animals to Humans." *Journal of the American Medical Association* 296 (2006): 1731-1732.

Hare, R. M. *Moral Thinking: Its Levels, Method and Point*. Oxford, UK: Clarendon Press, 1981.

Heap, Shaun Hargreaves, *et al. The Theory of Choice: A Critical Guide*. Cambridge, MA: Blackwell, 1992.

Howard-Snyder, Frances, and Alastair Norcross. "A Consequentialist Case for Rejecting the Right." *Journal of Philosophical Research* 18 (1993): 109-25.

Jaschik, Scott. *Did Peter Singer Back Animal Research?* 2006. <http://www.insidehighered.com/news/2006/12/04/singer> (14 February 2007).

LaFollette, Hugh, and Niall Shanks. *Brute Science*. London, UK: Routledge, 1997.

Luce, R. Duncan, and Howard Raiffa. *Games and Decisions: Introduction and Critical Survey*.

New York, NY: Dover Publications, 1985[1957].
Mill, John Stuart. *Principles of Political Economy with Some of Their Applications to Social Philosophy*. In *Collected Works Of John Stuart Mill*, *Vol. III,* edited by J. M. Robson. Toronto O.N., Canada: University of Toronto Press, 1965.
_____. *Utilitarianism*. In *Mill's Utilitarianism: Text and Criticism*, edited by James M. Smith and Ernest Sosa. Belmont, CA: Wadsworth, 1969[1861].
_____. *Whewell on Moral Philosophy*. In *Collected Works of John Stuart Mill*, *Vol. X: Essays on Ethics, Religion and Society*, edited by J. M. Robson. Toronto O.N., Canada: University of Toronto Press, 1969[1852].
Mukerjee, Madhusree. "Trends in Animal Research." *Scientific American* 276 (1997): 86-93, 1997.
Norcross, Alastair. "Good and Bad Actions." *The Philosophical Review* 106 (1997): 1-34.
Norcross, Alastair. *Reasons without Demands: Rethinking Rightness.* In *Contemporary Debates in Moral Theory*, edited by James Dreier. Malden, MA: Blackwell Publishing, 2006.
Nozick, Robert. "Interpersonal Utility Theory." *Social Choice and Welfare* 2 (1985): 161-179.
Perel, P., *et al.* "Comparison of treatment effects between animal experiments and clinical trials: systematic review." *BMJ*, doi:10.1136/bmj.39048.407928.BE (2006): 1-6.
Perry, P., *et al. The Ethics of Research Involving Animals*. 2005. <http://www.nuffieldbioethics.org/fileLibrary/pdf/RIA_Report_FINAL-opt.pdf> (20 January 2007).
Plous, Scott, and Harold Herzog. "Reliability of Protocol Reviews for Animal Research." *Science* 293 (2001): 608-609.
Pluhar, Evelyn B. *Beyond Prejudice: The Moral Significance of Human and Nonhuman Animals*. Durham, NC: Duke University Press, 1995.
Pound, P., *et al.* "Where is the evidence that animal research benefits humans?" *BMJ* 328 (2004): 514-517.
Railton, Peter. "Facts and Values." *Philosophical Topics* 24 (1986): 5-31.
_____. "Moral Realism." *Philosophical Review* 95 (1986): 163-207.
Rennie, John. "The Animal Question." *Scientific American* 276 (1997): 4.
Sidgwick, Henry. *The Methods of Ethics (7th ed.).* Indianapolis, IN: Hackett Publishing Company, 1981[1907].
Singer, Peter. *Animal Liberation: A New Ethics for Our Treatment of Animals*. New York, NY: Avon Books, 1975.
_____. "Utilitarianism and Vegetarianism." *Philosophy and Public Affairs* 9 (1980): 325-337.
_____. Setting limits on animal testing. *Times Online*. 2006. <http://www.timesonline.co.uk/tol/comment/article658128.ece> (14 February 2007).
Slote, Michael. *Common-Sense Morality and Consequentialism*. Boston, MA: Routledge and Kegan Paul, 1985.
Sorabji, Richard. *Animal Minds and Human Morals: The Origins of the Western Debate*. Ithaca, NY: Cornell University Press, 1993.
Stephens, Martin L. *A Current View of Vivisection: Animal Research in America.* In *A Primer on Animal Rights*, edited by K. W. Stallwood. New York, NY: Lantern Books, 2002.