taking into account both human and animal concerns, with serious attempts to reduce significantly the harm done to the animals. In any case, these are the most ethically agonizing uses, requiring the most compelling justifications, practical (e.g. the relief of a great human harm) or theoretical. In actual practice, there might be little difference between these two ethical positions.

On one point we were in unambiguous agreement. This was the dismay over the current, sometimes violent contention and level of public debate that surrounds the use of animals in science. Not only is there barbarism and intransigent, ill-suited ethical dogmatism. More importantly there is a blindness to subtle issues that we are only beginning to recognize and that urgently need exploration, such as the intimate and concrete interrelations of individuals, communities, and ecosystems that involve humans, animals, and other forms of organic being alike. We desperately need to think systematically about our ethical responsibilities to these various interconnected forms and kinds of organic, biological life. But this is impossible if we shortsightedly pit humans against animals against the environment. Such ethical provincialism is a curse we no longer can afford. **—Strachan Donnelley**

Section I. Ethical Theory and the Moral Status of Animals

Members of an institutional animal care and use committee (IACUC) spend much of their time inspecting laboratories and care facilities, and studying regulations. Yet one of their most important duties is to grapple with ethical problems: whether specific experiments involving animals are morally justifiable and whether they should be permitted. Yet this is often also the duty with which IACUC members feel most uncomfortable.

Understandably, perplexed decisionmakers might turn at this point to moral philosophy for some guidance. They will be sadly disappointed if they hope to find an adequate and generally accepted moral theory that will generate unassailable, defensible solutions. For every ethical theory advanced in philosophic discussion, several oppose it, and consensus among moral philosophers is not to be found. Nevertheless, moral philosophy can still make an important contribution to the debate about the moral status of animals and their use in biomedical research by analyzing relevant intuitions and arguments. Our intuitions, specific decisions, and even feelings of conflict are implicitly shaped by and reflected in the ethical theories that historically have engaged the attention of moral philosophers. We can better understand and evaluate those intuitions and decisions if we understand more clearly the theoretical commitments that underlie them. Even our disagreements about specific issues often implicitly reflect deeper conflicts between moral theories. Understanding the nature of these conflicts is a necessary first step toward resolution.

Thus, the following survey of ethical theories is intended to serve as a sort of toolbox for IACUC members to help them refine, evaluate, and justify their decisions, and perhaps even to resolve or at least better understand conflicting intuitions. We shall pay particular attention to ways in which these various ethical theories fit together, and to how tensions between them reflect our conflicting intuitions about specific cases. This approach is grounded in a fact about both real-life ethical dilemmas and abstract debates in moral philosophy: a satisfactory answer is most likely to be found in a judicious combination of theoretical commitments. However, a plea for tolerance and compromise should not be confused with accepting inconsistency, carelessness, or arbitrary switching from one theoretical ground to another—whether for the sake of convenience or to preserve intuitions and prejudices.

Utilitarianism

The most important version of *teleological* or *consequentialist* theories—theories that focus on the consequences of actions—is "utilitarianism," which was first explicitly articulated under that label by Jeremy Bentham (1748-1832), and developed more fully by John Stuart Mill (1806-1873).[4,22] The basic theory that grew out of their philosophy, now generally known as *classical utilitarianism,* is essentially this: In deciding whether an action is morally right, we sum up the total amount of good the action will bring about, and weigh that against the total amount of harm that will be caused. An action is right if and only if it brings about a better balance of good consequences over harm than any alternative action. Bentham explicitly identifies the "good" in question as pleasure or happiness; pain, suffering, or the diminution of happiness constitutes "harm."

Classical utilitarianism appears to fit well with many of our ordinary decisions, especially about the use of animals in research: Probably the most common justification is that the benefits gained from such research (e.g., curing debilitating illness) far outweigh the suffering of experimen-

tal animals. However, this argument must be examined more closely. Inasmuch as animals are capable of feeling pleasure and pain we must include these goods and harms in our utilitarian calculation of whether an action is morally right. Bentham recognized this explicitly,[4] and a contemporary utilitarian, Peter Singer, has applied such calculations to animal research and concluded the vast majority is immoral.[40]

This general argument has two prongs. First, much research produces trivial, if any, benefit, and hence has little to weigh against the suffering caused to the animals used. To complicate matters, utilitarian calculations usually require us to factor in the *probability* of an expected benefit or harm, and since no experiment is certain to produce a hoped-for benefit, the "expected utility" will be reduced still further. Second, even if a particular research project promises significant gains, there may be alternative ways of alleviating suffering with less harm done. For example, instead of using funds for developing a new technique for open-heart surgery, we could do more good by using that money to convince people to change their diet and exercise patterns. Thus, classical utilitarianism issues a challenge to re-evaluate our almost casual acceptance of current levels of use of animals in research.

Despite its obvious appeal, classical utilitarianism has been subject to important criticisms. First, the method of summing up or averaging the good or harm for all concerned is often seen as failing to respect those affected as separate and distinct individuals. For instance, utilitarian theories seem to condone secretly killing someone who contributes no good to society, friends, or family in order to transplant her organs into four or five worthy but dying people. Moreover, hedonism seems too narrow as a full account of "the good"; after all, many people freely forgo pleasure to obtain other things they deem worthy of pursuit.

In response, contemporary utilitarians have offered two important modifications to classical utilitarianism. First, many advocate some version of *rule utilitarianism*. That is, instead of evaluating actions one at a time, we should identify and follow rules which, if adopted, would maximize good consequences. Thus, an action may be wrong, even though it produces more benefit than harm, if it violates a rule that would, in general, yield better overall consequences. Second, many utilitarians agree that there are many goods in addition to pleasure that ought to be maximized, including individual autonomy, the satisfaction of preferences and interests. These accounts are often referred to as *pluralistic* versions of utilitarianism.

Unless one is prepared to argue that animals simply lack the capacity for morally relevant interests and desires,[18] any consistent form of utilitarianism that includes pleasures, suffering, or interests in its account of "good" must accord sentient animals a significant moral status. An act may therefore be wrong solely because of its direct effect on animals, even if it causes no harm to humans. Torturing kittens may or may not foster a tendency to be brutal to humans, but according to utilitarianism, it is wrong essentially because it harms the kitten. Although this view entails that any being capable of suffering or having interests is morally entitled to *equal* consideration, it does not follow that animals have a moral status equal to that of human beings, or that the life of an animal and that of a human are of equal worth.[40] Many values, such as choosing and carrying out a long-range life plan, moral autonomy (the ability to function as a rational moral agent), and complex personal and communal relationships, are relevant exclusively or to a much greater extent for normal human beings. One is justified in giving "preference" to human beings insofar as doing so maximizes all good consequences, including those that have special value for human beings. In this way, utilitarianism does justice to the conviction that animals have *some* moral status—that benefits or harms to animals must figure *directly* into our judgment that an action is right or wrong—but also provides justification for asserting that the life of a human being may have more value than the life of one or many mice.

A pluralistic version of utilitarianism may also justify the intuition that, other things being equal, it is better to use a planarian than a mouse, and better to use a rat than a chimpanzee. A mouse is not only capable of feeling pleasures and pain, but also has other interests and preferences that can be satisfied or frustrated. A chimp has even more such interests, a richer network of personal relationships, and a greater sense of the future. For this reason, significant interference in the life of a chimpanzee is likely to produce more bad consequences than the same sort of treatment of a mouse.

Concern for moral autonomy and other "higher" goods may supplement but nonetheless cannot supplant the demand to maximize happiness and the satisfaction of basic preferences or interests. Once we include the suffering and frustration of interests or preferences of research animals in our utilitarian calculation, we will be driven to the conclusion that a sizable proportion of current research is morally unjustifiable. Rule utilitarianism, no less than classical utilitarianism, demands a strict, fair, and complete "cost-benefit" analysis of the sort that many IACUCs today—in their insistence that considerations of scientific merit are outside their province—in fact abjure.

Deontology and Rights

The second important category of ethical theories is known as *deontology*. According to a standard definition, a theory is deontological if and only if some acts are judged right [wrong] even though their consequences are, on balance, bad [good]. That is, a deontological theory rejects, at least in part, a utilitarian approach to ethics. The qualification "at least in part" acknowledges that many deontologists (notably, for our purposes, Tom Regan) distinguish "moderate" from "extreme" versions of deontological theories. "Moderate" theories allow consequentialist considerations to play a substantial role in moral deliberation, while still maintaining that "rights are more basic than utility and independent of it."[35]

The rejection of consequentialism

entails that we should not make moral judgments simply by asking whether the consequences of an action would result in the world's being "a better place" (happier, more pleasant experiences, more satisfaction, etc.). Instead, most deontological theories offer a concept of right (as in "right action," not necessarily "a right to x") that is independent of or logically prior to any specification of "the good." This emphasis on a special class of irreducible moral properties or "right-making characteristics"—what makes an action right—of actions is especially common in twentieth century revivals of deontological theories.[34,36]

The distinction between "right" and "good" suggests that the two concepts are to some extent incommensurable, even though a moderate theory should have some way of balancing the two. That is, in a strict deontological theory one cannot justify an action that would otherwise be wrong simply on the ground that it produces more good consequences than harmful ones. The rightness or wrongness of an action forms a barrier that cannot be breached by an appeal to consequences. If torture is wrong, we cannot justify torturing a terrorist merely because we can achieve considerable good by extracting information in this manner. For this reason, deontological theories often take the form of a *rights-based* account of moral obligation. More specifically, serious talk about "animal rights" is more likely to find conceptual justification in deontology.

Deontological theories differ widely on the issue of the moral status of animals and their use in research. Immanuel Kant (1724-1804), who formulated the classic version of deontology, claimed that we have no direct duties to animals, that they have no moral status, and would therefore condone any research that might benefit human moral agents. By contrast, Tom Regan, a contemporary deontologist, argues that animals do have significant moral status, and that almost all animal research is immoral. This difference can be traced back to a disagreement about "right-making characteristics" and the more basic concept of inherent value.

Deontologists can give substance to their disagreement with consequentialism by insisting that some things have value in virtue of the sorts of things they are, not simply in virtue of their contribution to the total amount of good in the world. The locus of inherent value is reflected in a concern for individuals as individuals of a certain sort. To use a metaphor that Regan often invokes, the individual is not merely a "receptacle" for some impersonal good; he/she/it has moral standing simply in virtue of the sort of individual he/she/it is. Consequently, we cannot merely sum up the total amount of good contained in all the "receptacles," or transfer good from one receptacle to another, without considering the inherent value of the receptacle itself.[35]

Animals, then, have moral status insofar as they have inherent value, that characteristic by virtue of which they are worthy of moral respect. Unless we follow Kant in attributing inherent value only to moral agents, in which case we will have problems justifying our respect for infants, severely retarded humans, or other members of our own species incapable of acting as a moral agent, it is reasonable to expect that at least some of the properties in which inherent value is grounded will be found to some extent in non-human animals. It is hard to move beyond this point, since the basis of inherent value is the subject of much dispute. For example, Kant took moral autonomy as the basis for inherent value, while Regan argues that we ought to recognize inherent value of all "subjects of a life."[35,39] These are by no means the only alternatives; whenever our intuitions prod us into an attitude of respect, awe, or wonder at something as it exists for its own sake, invoking a special moral commitment to its preservation, we may well be drawing on a deontological concern for inherent value.

Philosophic discussion of inherent value is necessarily inconclusive; the concept nonetheless still helps to account for some of our intuitions about the use of animals in research, and to challenge others. As we have just seen, inherent value can explain why we ought to respect animals as individuals with lives of their own. The deontological emphasis on respect and justice for the individual often underlies the judgment that it is wrong to be callous or casual about the use of animals even where those attitudes don't result in actual harm.

The reaction of most people to the videotapes of research on unanesthetized primates at the University of Pennsylvania trauma lab exemplifies this perspective; the callous and casual attitude of the workers may have been at least as objectionable as the actual abuse. There is simply something morally wrong about laughing and joking about the condition of a helpless subject, whether or not any additional pain results. While there are serious difficulties in articulating what the moral obligation to respect other individuals requires and what sort of beings are entitled to respect, we still recognize that our moral theories ought to incorporate this requirement.

The deontological notion of inherent value and respect also affects our judgments about research in which animals are killed without experiencing any pain or distress. Utilitarians (although they have often struggled with this issue) have great difficulty explaining what might be morally wrong about killing an animal painlessly, especially if we have good reason to think that it is incapable of sophisticated concepts and desires such as the desire to live long enough to carry out a long-range plan.[41] Nonetheless, many people are convinced that an experiment in which five animals are painlessly sacrificed is preferable to one in which ten are used to achieve the same result, and that any interference with the life of an animal, no matter how painless and innocuous, should require some moral justification. No animal should be killed wantonly.

Despite the fact that deontology may attribute inherent value to at least some animals (Regan, for example, takes mammals over one year of age as his paradigm), deontology's particular emphasis on specifically *moral* values suggests that moral autonomy will carry special weight. Thus, deontology, like utilitarianism, leaves room for arguments about the special status of human beings, and for

appealing to a hierarchy according to which it is better to use a more "primitive" animal rather than a more complex one.

Unlike utilitarianism, however, deontology produces confused results when addressing the issue of "higher" versus "lower" animals. Inherent worth is often taken to be an all-or-nothing characteristic, although it is not clear that this is a logically necessary feature of deontological theories. Given the lack of consensus on this point, the best we can say is that our belief that we should use "lower" animals rather than "higher" ones wherever possible is more congenial to utilitarian justification. This is yet another illustration that no single, "pure" moral theory has emerged as completely adequate for moral decisionmaking, and that our specific judgments will draw on a combination of approaches.

Contracts, Kinship, and Organic Nature

Philosophers often discuss moral theories that do not fit neatly into the utilitarianism/deontology dichotomy. While these theories are sometimes treated as alternatives to utilitarian and deontological approaches, we think it is more useful to see them as attempts to identify some of the values that, while compatible with either utilitarianism or deontology, are often neglected. We will here focus on three topics most directly relevant to the moral status of animals: contracts, kinship and community, and organic unity.

Contractarian theories of ethics see moral obligations as growing out of mutual agreements between moral agents: I will refrain from doing x to you if you, in return, agree not to do x to me. Since most moral constraints are never formulated so explicitly, contractarians typically present them as implicit agreements or even as convenient fictions—what a rational person would agree to if she were asked to enter into a new society of peers. Some philosophers (e.g., Thomas Hobbes) have argued that contracts would be designed to maximize benefits for all concerned, and so would be compatible with utilitarianism. Others (e.g., John Rawls) emphasize the fact that moral agents would, first and foremost, take care to secure the liberty necessary to function autonomously, and ally themselves with deontological approaches.

Contract theories are by nature limited in scope. Only moral agents can enter into a contract; thus, any obligations we might have to infants, severely retarded people, and other humans who cannot function with the required degree of moral autonomy, will not follow directly from the basic premise of such theories. On the other hand, they may explain some special obligations to human society; for example, a scientist's obligation to "repay" human society for her education and research support by emphasizing research that will most benefit that group.

Kinship and the related concept of social group or community often surfaces in discussion of the use of animals in research.[21] It is often asserted, for example, that we have stronger obligations to our own children and other family members than to strangers; this may be generalized to include one's community as a whole. Usually, this entails favoring one's own kin and conspecifics (one's species) over the lives of other animals. The same attitude can also be seen in the special obligations we seem to feel towards our own pets, domestic animals who have become part of our social circle. Some philosophers (particularly utilitarians) have argued that this preference is irrational, that a good moral theory must be impartial and completely fair. As our objective here is not to resolve philosophical disputes but to delineate connections between philosophical concepts and specific decisions, it may be sufficient to articulate the point clearly, and to see how it fits with the two main types of theories already identified.

Utilitarianism is often criticized for being impersonal, treating people as completely interchangeable, but even utilitarian theories can accommodate the special demands of kinship by arguing that good consequences will be maximized if agents concentrate their efforts to help those they know and understand best—family, friends, fellow citizens, and even conspecifics. Within a deontological framework, certain relationships may in and of themselves generate specific personal duties. Indeed, it is often deemed an argument in favor of deontology that it can account for the fact that our obligations to family members often count for more than our duties to anonymous strangers.

Talk of "organic unity" represents a unique attempt to combine ethics and metaphysics.[14] The claim is that organic life exhibits a kind of individual unity, dynamism, purpose, and "neediness" that gives it inherent moral value. Organisms and communities of organisms create, instantiate, or define ends in themselves, simply in virtue of leading a purposeful existence. Organic unity similarly presents a value that might be incorporated into either a utilitarian or a deontological theory: the goals exhibited by an organic unity might be taken as good consequences to be maximized, or as states inherently worthy of a special moral respect.

This value system, perhaps more than any other, captures many features of our basic attitudes toward animals. It recognizes life and the striving purposiveness that accompanies life as morally significant. Moreover, since complexity adds richness and value, more complex organisms are more valuable than simple ones. This value extends beyond the individual to the entire community in which the organism operates; here too, complexity and richness are characteristics to respect and protect. Thus an endangered species, or one that is part of an intricate and delicate ecosystem, gains extra value from that relationship, and therefore deserves more care and respect than, for example, a laboratory mouse.

Value systems such as these seem to capture the essence of certain very powerful intuitions we have about how we ought and ought not treat animals, and as such, must be acknowledged and investigated further.

Moral Reasoning

What are we to make of the preceding smorgasbord of alternative ethical theories and value systems?

We have tried to show how moral philosophy cannot give us a neat formula or algorithm with which to generate safe and defensible solutions but can still help us understand, evaluate, and occasionally change the decisions we do make. By focusing on the ethical theories that underlie our moral decisions, we are forced to confront considerations of consistency and justification that otherwise remain hidden. We cannot, for example, argue that an experiment that should not be done on a dog is permissible when pigs are substituted if we are not prepared to identify a morally relevant characteristic of dogs (or pigs) that would justify the different evaluations.

When we move from abstract theories to specific issues, we find important areas of agreement. The strongest and most defensible ethical theories all entail that animals—at least those which are sentient (capable of feeling pleasure and pain) or "autonomous" in the wide sense of having preferences and the ability to pursue them—have a significant moral status. They cannot be treated as mere objects, and the effect of an experiment on them must be considered when we decide whether the experiment is ethically justifiable. However, the values with which ethics is concerned are numerous and varied, which suggests that moral status is not an all-or-nothing characteristic. That is, the value systems most adequate for real dilemmas will recognize happiness, absence of pain, autonomy, experiential and cognitive richness, and many other factors as relevant to moral judgments. Since these properties are more relevant to the lives of some animals than to others, the result is likely to be something like a "moral hierarchy" of concern (although the relative weighting within the hierarchy may shift from case to case, as different concerns become relevant).

In adopting this approach, we have discovered that different theories and value systems are needed to do justice to different aspects of the decision-making process, and to capture the conflicts that arise in hard cases. This reinforces the message conveyed by the history of moral philosophy: An adequate ethical theory is most probably going to contain elements of both utilitarianism and deontology, and will involve a pluralistic value theory.

None of this is a license for promiscuity. One cannot justifiably select whatever moral theory happens to generate the conclusion one would like to reach. Legitimate areas of agreement, overlap, and compatibility must be carefully articulated and analyzed; our foundational theory may be eclectic, but cannot properly contain inconsistencies or arbitrary choices. To be sure, this is a difficult enterprise, but one that should result in a clearer understanding of our obligations regarding the use of animals in biomedical research.

—Lilly-Marlene Russow, in collaboration with K. Danner Clouser, David DeGrazia, and James Stephens

Section II. Animals in Science: The Justification Issue

At the heart of the wide-ranging and seemingly unending controversy over the use of animals in biomedical science, whether in basic or applied research, toxicity testing, drug production, or education, is one burning question: Are we humans ethically *justified* in such a use of animals, in general and in particular cases? How we individually or collectively answer this question no doubt finally depends on our moral worldview, including the judgment of the relative moral status of human beings and animals, discussed in the preceding section. However, the very fact that controversy continues to rage among the more extreme animal liberationists/rightists and human welfarists/scientific progressivists, with the "troubled middle" caught in-between, provides evidence for a plurality of "ultimate" or primary ethical values, goods, and responsibilities that are not easily commensurable.

In addressing the justification controversy, it is important to see precisely what is at issue. The central factors leading to conflict are the nature and methodological requirements of scientific inquiry, the human values of science, and the complex values accruing to animals and organic life, including the ethical responsibilities that these values engender. We are caught in a cross current of seemingly ultimate ethical values and obligations from which we do not easily escape.

The Nature of Scientific Inquiry

Certain general features of biomedical science (understood here as the scientific study of organic life for both theoretical and practical reasons) require notice. Whatever be the most adequate theory of the scientific method—the so-called hypothetico-deductive method with its infusion of rational imagination—modern biomedical science is characterized by controlled experimentation for validation of scientific ideas, knowledge, and techniques. Moreover, biomedicine essentially involves an interdependence of various scientific disciplines and areas of inquiry for generations of important scientific knowledge and practical techniques.[11,19] This interdependence extends to the various "models" used in experimentation, whether mathematical, cellular, animal or human, *in vivo* or *in vitro*.[19,45] Animal exper-