A Constrained-Utility Alternative to Animal Rights

JOSHUA FRANK

Foundation for Interdisciplinary Research and Education Promoting Animal Welfare (FIREPAW)
221 Lark Street
Albany, New York 12210, USA
Email: firepaw@earthlink.net

ABSTRACT

Numerous approaches have been taken in an effort to find a non-anthropocentric ethic that will lead to greater consideration of animals. Most of the recent approaches in this area have been rights-based. It is argued here that a rights-based approach alone fails both theoretically and in practical applications. It is shown that in theory these approaches can lead to unsound conclusions and cannot handle uncertainty. In addition, in practice the rules of the rights-based approaches will often be violated. A utility approach with unequal weighting for different species subject to certain rights or obligations is proposed as an alternative. This approach is intended to be operational rather than purely theoretical and therefore would be based on a negotiated consensus rather than a priori theory.

KEY WORDS

Anthropocentrism, speciesism, utility, rights, moral status

I. THE RIGHTS-BASED PERSPECTIVE

There is a growing sentiment that animals should be given greater ethical consideration. However, there is much less consensus in defining exactly why and how animals should be valued.

The Case for Animal Rights by Regan (1983) presents a rights-based argument for granting animals greater ethical consideration. Specifically, Regan argues that animals possess a moral right to be treated with respect. He argues at length that the position taken by Singer (1975) and others who base animal interests on utility considerations is flawed. One argument Regan advances

Environmental Values **11** (2002): 49–62

against utilitarianism is that it makes killing too easy. If the utility gained by such an act outweighs the utility lost, then death even of an innocent person or animal is called for according to a purely utilitarian perspective. This may seem like a difficult outcome to accept, but given the context (i.e. the value of utility gained is more than the value of an entire life), why is it not the correct choice?

Before dismissing utility-based values in the light of this argument, we need to examine in more depth why people find this situation objectionable. Much of our distaste with making such choices may be based on practical considerations rather than purely theoretical considerations. There are some very good practical reasons for wanting to avoid killing even if it will cause a net utility gain. First, it fails to consider uncertainty. Future outcomes are virtually never known with certainty. Therefore, killing now for a supposed benefit later may be difficult to justify as a practical matter. Furthermore, decisions always involve a decision-maker who is a human being subject to potential error and corruption. Therefore, putting the ability to make life and death decisions in anyone's hands may be unacceptable because of the possibility of mistakes or abuse of the authority (some opponents of the death penalty use these first two arguments – even if killing is considered justified if all facts are known, there is human error and inherent uncertainty that can lead to innocent people being sentenced to death). And finally, there is the 'slippery slope' argument: allowing the life or death decision once may lead to such decisions being made in more questionable circumstances. Therefore, if we look deeply at our opposition to killing for net utility gains, we may often find that it is based more on practical than theoretical considerations.

If a benevolent, infallible dictator were to exist who had no interest in the consequence of decisions and acted in the interest of all beings with moral status, the idea of sacrificing a life for a net utility gain for all might be more acceptable in some cases.

II. PROBLEMS WITH A PURELY RIGHTS-BASED PERSPECTIVE

A purely rights-based perspective has its own difficulties. Though a rights-based system such as the one proposed by Regan seems to avoid any sort of hedonic calculus on the surface, when examined closely this is not really the case. When challenged, these systems must resort to some form of utilitarian calculus to avoid becoming incoherent or inconsistent.

Regan seeks to preserve the position that no possessor of rights should receive a great harm to give many individuals a smaller benefit, while at the same time maintaining that harming two people is worse than harming one if the level of harm is equivalent for all. He therefore puts forward the 'worse-off' principle. Based on the idea of the right to respect, Regan posits that in any situation where a choice must be made between two actions that will both cause harm, the

greatest harm to any single individual should always be avoided. Only when all harms are equal should the number of individuals receiving the harm be taken into consideration. But this position does not hold up well under scrutiny.

First, consider that 'harms' in reality exist on a continuum. Therefore some arbitrary level of significance must be used for determining when harms are 'equal' or when a harm is worse. Regan uses an example of saving thirty people in one part of a mine versus saving one in another part. The thirty should be saved he says, since the harm is equal to all. However, let us say that we know that the one person will slowly suffocate while the others will die quickly in a cavern collapse. Is the one person now suffering a greater harm? And if so, would we really prefer to save the one to the thirty? To use an analogous situation, if breaking two arms and one finger is worse than just having two arms broken, would we really prefer to break two arms of a thousand people to breaking two arms and one finger of one person? This is the conclusion that Regan's logic would lead to.

Introducing uncertainty makes the situation more difficult and unacceptable. Let us take the same mining scenario, now say that the one person has a 99% of dying, while the 30 people have a 98% chance of dying. Do we save the one person based on a 99% chance of death being a greater harm than a 98% chance of death? This does not seem like the acceptable choice, especially if we increase the 30 people to 30,000 people and make the percentage difference 99.01% vs. 99% or any arbitrarily close but distinct percentages. Or if Regan would argue that 99% of dying is not a greater harm than 98% of harm (since in either case death is still the harm), then this must hold for any percentages. In other words, if a 99% chance of death is the same harm as 98% chance of death, there is no reason to hold that a 1% chance of death is worse than a 99% chance of death. The only solution that is consistent with our intuition in this case is to resort to a pseudo-utility theory arithmetic by calculating expected values or some other risk-weighted function.

The situation becomes even more complicated if we now make the more realistic assumption of a continuous probability distribution. For example, Regan acknowledges that the actual sensitivity of people to harm such as pain may vary. He simply assumes the sensitivity is equal in all cases. But what happens if we take our knowledge and construct a probability function of harm from a stimulus for that individual? What is the 'greatest harm'? Do we use a minimax principle and minimise the greatest harm? This may not even be possible, for a continuous normal distribution for example, and it may not make logical sense if we make the probability of the highest damage scenario arbitrarily low. Once again, we are forced into the only sensible solution being a probability-weighted utility calculation.

Regan also presents a lifeboat scenario in his book. Four people and one dog are in a lifeboat, and one must be tossed overboard for the others to survive. Regan essentially forces his desired solution of choosing the dog by stating that

though all deserve equal respect, the dog's life is somehow less valuable. Regan reasons,

> all on board have equal inherent value and an equal prima facie right not to be harmed. Now, the harm that death is, is a function of the opportunities for satisfaction it forecloses, and no reasonable person would deny that the death of any of the four humans would be a greater prima facie loss, and thus a greater prima facie harm, than would be true in the case of the dog.

Since this is based on the respect principle, Regan would reach the same conclusion regardless of whether the choice is one dog or a million dogs.

There is no real basis for stating that killing a human is a greater harm using Regan's logic since there is also little basis for a human to judge a dog's 'opportunities'. The 'opportunities for satisfaction' justification lacks substance and can lead to conclusions many would find unsatisfactory, for if it is opportunities that are important, the answer should change if one of the people is terminally ill and only will live one year while the dog is expected to live ten more years. Or what if the person 'lacks opportunity' for some other reason such as a handicap? Does that now mean he or she should be thrown overboard? According to this logic, we should throw ten thousand such people overboard rather than one person with slightly greater 'opportunity'. Indeed, since no two people ever have the same level of opportunity, the concept of equal harm becomes meaningless. The situation once again becomes even more complicated if the probability of death is introduced, and even more complicated if uncertainty in opportunities is considered.

III. THE DISCRETE CATEGORY PROBLEM

There are at least two major issues here. The first one could be called the 'discrete category problem' and stems from the use of discrete categories for harms (and benefits) that are in reality continuous. Due to this problem, a rights-based system fails to give guidance or leads to unsound conclusions when faced with complex ethical dilemmas, particular when information is quantified. This problem is not unique to Regan's particular system of rights. The key reason that a rights-based system is often advocated over a utilitarian framework is that rights supposedly avoids distasteful trade-offs such as taking a life for a small gain in utility for a large group of people. In particular, most systems of rights dealing with animals are engineered to allow important needs such as survival or avoiding severe suffering to a sentient animal always to take priority over minor pleasures or benefits for human beings. At the same time these rights-based systems often allow human interests to take priority when the level of the need is identical (such as death to a human versus death to an animal).

In a pure utilitarian framework where some social utility formula is simply maximised, animal interests can always prevail over even the survival interests of a human being, provided animals are given some weight in the calculation and enough animals are affected. Conversely, a minor human interest can prevail over a survival animal interest if a large enough relative weight is given to human interests or if the number of humans affected is sufficiently large.

A system of rights avoids these trade-offs by avoiding trade-offs altogether. This is done by creating absolute categories of value rather than a continuum. For example, VanDeVeer (1979) divides interests into three categories: 1) basic interests necessary for survival, 2) non-basic but serious interests, and 3) non-basic, peripheral interests. Taylor (1986) similarly divides interests into central and peripheral interests. Regan, as previously mentioned, seems to allow for a continuum by using the 'equal harm' concept, but Regan's concept is not a continuum because categories are treated distinctly. No matter how large or fine the categories are, the problem remains the same as long as ethical treatment is discrete based on category rather than mathematically weighted based on a continuum. If categories are assumed to be very small, then a slightly smaller harm to a very large number of beings having moral status is preferable to a slightly larger harm to a single being. If decision rules are based on very large, general categories (such as basic, serious, and peripheral interests) then there is no distinction between interests within a category. Moderate but serious suffering must be on par in our value system with intense but non-fatal suffering. If the two are not treated as equal, then the same question can be asked going halfway between the new subcategories until either two very distinct harms must be treated equally or the categories become so close together that the dilemma discussed earlier for finely defined categories occurs.

IV. THE UNCERTAINTY PROBLEM

The second issue discussed is the uncertainty problem. Of course, no ethical discussion can account for every complex situation and it is unfair to criticise a value system for failing to state explicitly how certain unforeseen questions would be handled. But this is not simply a matter of bringing in an excluded added layer of complexity. The issue is that there is no coherent way to include uncertainty in a rights-based system with discrete categories of harm without leading to unacceptable results.

There are four ways to handle a chance of harm (or benefit). It can be treated as the greatest possible harm, it can be treated as the least possible harm, the probability can be arbitrarily divided into categories such as 'likely' or 'very likely', or a probability-weighted function such as expected value can be used. Treating the harm as the greatest harm or treating it as the least harm ignores the

probability of harm and leads to undesirable results. Categories of probability are not only arbitrary but also lead to the same difficulty previously discussed for categories of harm in general. The most sensible solution is an expected value or other probability function. This works for a value system based on utility. But using expected value for categories such as basic versus peripheral interests or for deciding what is a 'greater harm' is a meaningless exercise.

It is important to note that this is not an exceptional case. Uncertainty is the norm in real ethical dilemmas pertaining to human versus animal interests. The clearest example of this is animal research. Although it is often described otherwise by proponents of animal research, the fact of the matter is that any animal research has at best a chance of benefit for human life. Comparing certain animal suffering and death to possible medical human benefits is not even possible without considering uncertainty (unless of course the exact same level of possible benefits can be achieved without animal suffering in which case the solution is clear). Most dilemmas involving wild animals also involve uncertainty. Development involving habitat destruction normally includes a chance of harm to certain animals. Human activity that poses hazards to animals (such as possible oil spills, possible chemical spills, or pollution with unknown effects on all species) is by definition uncertain in impact.

V. RIGHTS IN PRACTICE

In practice, most systems of rights often provide surprisingly little conclusive guidance when faced with real decisions. Take the case of a building a road in a non-urban area. As a consequence of constructing the road, there will be some mammals killed by traffic, sometimes with considerable suffering. Presuming other routes already exist, the only advantage to humans of the construction is decreased travel time. Most rights-based systems that seek to give animals greater consideration than they currently receive would suggest that the survival needs of the animals outweigh a peripheral human need. Yet if this logic is applied to all roads, only a primitive and time-consuming network would exist, which many would find to be an unsatisfactory solution.

Another completely different example where rights yields little insight is the overpopulation of unwanted dogs and cats. This is a case where the analysis of rights versus utility for purely human questions may differ when animals are included since our role as steward comes into play. It is certainly human action that has created the problem, however the actions that cause overpopulation are mostly indirect. Therefore, positing a right to life or even a right to a home (since these animals have been bred to be domesticated) yield no guidance in terms of what to do about the problem of existing animals with no home or how to address the problem.

One possible direction to take here is to use a system based on obligations rather than rights as suggested by O'Neill (1997). However, obligations are little better than rights in both cases. It is hard to say what obligations would lead to a proper balance of roads versus wildlife protection nor what obligations that would be considered acceptable could be morally required in the dog and cat overpopulation case.[1]

Another practical difficulty that arises particularly when animals are included in the picture is that in reality people trade between categories of harms and benefits. Despite distinctions theorists may make between interests that are basic, serious, and peripheral, the fact is that people will frequently trade some amount of basic interest for a peripheral interest (such as trading the possibility of death in a higher risk job for extra money that will be used mainly for luxury goods).[2] In a system of rights including only people, the fact that moral agents trade between categories according to their own preference may be of little consequence since everyone with moral status is also a moral agent who is free to trade as they please. But when a human's basic interest is favoured over the basic interest of an animal, and then that human trades that basic interest for a peripheral interest, in practice, a human's peripheral interest has been chosen over the basic interest of an animal. Therefore, if expressed human preferences violate the categories created by the theorists, as they often do in practice, the theoretical rules of what interests take priority will be violated in practice.

For example, let us assume that animal testing is banned for non-vital consumer products because it is decided that the right to survival of the animal outweighs the peripheral interest of humans in that product. But let us assume that there is a new anti-wrinkle cream with a new active ingredient that goes on the market even though it has not been tested for all health risks and there is a real possibility that the active ingredient could be carcinogenic. If certain people choose to consume that product even though there is a known health risk, then their expressed preference forces us to reframe the ethical debate. Because humans are definitely using the product this is no longer an issue of peripheral human interests. We now have a trade-off between a basic interest in the survival of animals with a basic interest in the health of humans.

We could also use a highway expansion example. If a major road expansion would violate the basic interests of wildlife but only serve the interest of increasing the safe and legal travel speed we have basic animal interests versus peripheral human interests in terms of reduced travel time. But many humans will most likely trade off their own survival interests for a peripheral interest by speeding and causing risk to life to reduce travel time. We now are forced into a comparison of a basic human interest in safety versus basic animal interests.

VI. UTILITY SUBJECT TO RIGHTS

I suggest the following rather simple social utility calculation be used as a basis for valuing animal and human interests:

$$\Sigma_t \Sigma_i \Sigma_s W_i P_{st} U_{its}$$

U is utility, P is the probability of a certain outcome, and W is the weight given to a particular individual's interests. The total is summed over all states of nature (s) all individuals (i) and all time periods (t). If time or the probabilities are treated as continuous, an integral could be used instead of a summation.

However, using a pure utility calculation for all situations can lead to conclusions many would find reprehensible, including Regan's example of taking a life for a minor benefit to many.[3] To counter this problem, certain basic conditions can be placed beyond the utility rules, and these conditions could be phrased in terms of rights or obligations although in some ways it does appear that obligations are advantageous. For example, there could be an obligation not to kill unless there is immediate threat to another life, an obligation not to intentionally cause suffering except in certain unusual and specific circumstances, and so on. These constraints (or rights/obligations) could also incorporate protecting non-sentient aspects of the environment. It should be noted that this is intended to be an operational approach rather than purely a theoretical exercise. Therefore, the constraints would be constructed procedurally through a negotiated consensus rather than based on a priori theory.

Our general rule then becomes:

Maximise: $\Sigma_t \Sigma_i \Sigma_s W_i P_{st} U_{its}$
Subject to: Right or obligation #1
Right or obligation #2
etc.

In decision sciences terminology this format is known as a constrained optimisation problem, and it allows us to create an explicit and operational rule that is flexible enough to incorporate the advantages of both rights and utility while avoiding most of the pitfalls of both. The 'subject to' language is merely a list of conditions that cannot be violated by the policy maker in attempting to maximise the equation.

An example of a law that utilises a constrained utility approach is the UK Animal Procedures Act (1986). According to this act, the decision to authorise a project must weigh 'the likely adverse effects on the animals concerned against the benefit likely to accrue as a results of the programme to be specified in the licence'. However, in addition to this weighting, the Act imposes certain constraints. For example, the certain animals including cats, dogs, and primates should only be used if no other species can be used. But more importantly, in any

event the animal should not be subject to severe pain, distress, or suffering. The Act also calls for the construction of an Animal Procedures Committee with diverse backgrounds to oversee the operation of the Act.

It should be noted that the utility calculation is intended as a social utility curve to be calculated in theory by an outside party for all individuals involved in a dilemma. The significance of this point is that it avoids the criticism that utility maximisation of an individual who may or may not exist is not meaningful. The equation is usable and meaningful in this sense: given the parameters, it can be calculated, therefore it is usable. And we know what the calculated number means, therefore it is meaningful, even if we cannot identify the specific individuals in a future time period. Specifically, the calculated number represents the summed utility expected by all members of society (including animals) combined regardless of who they are, weighted by the value placed on each individual.

A second obvious point is that the equation allows for different weighting being given to the harms and benefits received by individuals of different species. Although weighting could vary at the level of the individual, I believe a more reasonable methodology is to vary weighting at the species level. The idea of different weights for different species is not without controversy and to do this issue justice is beyond the scope of this analysis. Singer proposes that suffering is equal for all sentient beings. But others have made arguments for weighting species interests differently based on non-anthropocentric criteria. For example, Anderson (1993) justifies different intrinsic values for species based on the 'number of goods proper to the organism'; VanDeVeer bases preferential treatment on psychological capacity; Wetlesen (1999) argues for gradual moral status based on similarity to moral persons; and Lombardi (1983) argues for different inherent worth based on 'range of capacities'. However, the question of how and whether to give different weights to different species or individuals is too complicated an issue to address adequately in this discussion.[4] The important point is that the system described above is flexible enough to accommodate either equal weighting or a system of unequal weighting.

In addition to avoiding the pitfalls of rights-based systems, using an applied constrained-utility formula such as the one presented does have many advantages. One is that it leads to conciliation between opposing views. Rather than stating that animal rights always prevail in certain types of situations or that human interests always should prevail, it allows for compromise by allowing people from opposing camps to find ways to minimise the impact on the other side and by allowing the conclusion in each situation to depend on its own unique circumstances. It also explicitly takes into account all impacts and the precise degree of those impacts on sentient beings. In addition, although there would be considerable debate and controversy regarding the weights to give animals and the utilities of various outcomes, it is explicit and operational.

Of course, saying that this procedure is operational does not imply that it can be implemented lightly or without controversy. But the difficulty comes precisely from the fact that this equation is explicit in exactly what trade-offs are acceptable. Decisions must be made as to how many dogs' lives can be traded for a person's life or how much extreme pain is functionally equivalent to a single death. These may seem like difficult or even impossible trade-offs to make, but the fact is that policy-makers implicitly make such trade-offs every day. The only question is whether the values implicit in these decisions should be confronted head-on.

How do we then set the value for the equation parameters W_i, P_{st}, and U_{its}? Setting the probability of an outcome, P_{st}, is a technical matter that should be based on best scientific estimates. W_i (the value given to human and nonhuman life) and U_{its} (the utility value given to various outcomes) are more difficult since they are based on ethical judgments. For humans, in some situations this can be addressed using the economic concept of 'revealed preferences'. Using a revealed preference choice, relative weights would be based on the value people implicitly give to various options in the marketplace. However, this will not be useful for all ethical dilemmas, particularly when these dilemmas involve weighting human and nonhuman interests.

A more generalised methodology to use in such situations could be Social Judgment Theory as advanced by Hammond, Rohrbaugh, Mumpower and Adelman (1977). Hammond et al. applied their technique to the choice of bullets by the Denver police, a topic that had been subject to controversy. In their technique, scientific judgments were separated from value judgments in the decision process, with the former being made by experts and the latter being obtained from representatives of all parties concerned. In obtaining value judgments, the parties were asked to make hypothetical choices that were then translated into value weightings rather than being asked directly to come up with value weightings[5]. For the above equation, this technique can be applied to obtain values both for U_{its} and W_i.[6]

Just how this technique may be operationalised is best demonstrated by taking up a couple of the examples previously discussed.

As discussed in an Section V, a rights-based perspective offers little insight into how to address the problem of companion animal overpopulation which results in high rates of abandonment and death at shelters. On the other hand, the method outlined here can be successfully used to analyse this problem. The first step in the analysis process would be to define the relevant parties. Then responses would be gathered to assign weightings to outcomes relevant to this situation. In this particular case, utility weightings would be needed to compare possible trade-off such as the relative value of an additional dog placed in a good home versus one less dog euthanised, or the value of the life of a feral cat versus the value of the life of a cat in a home. Using this information, a utility function

can then be calculated and treatments can be recommended that maximise this function, subject to any right or duty-based constraint. A constraint in this case might be an obligation to kill no more animals in shelters than the current rate (or a more stringent version of this constraint could be that no companion animals would any longer be killed in shelters after a certain adjustment period). Public policy choices that can be analysed in this case include subsidised spey/neuter programmes, increased shelter space, and education programmes, to name just a few. The constrained optimisation function then gives a method to choose whether some, all, or none of these alternatives should be chosen.[7]

A similar process would be used in the animal research example. In this case, key factors to get weightings for include the relative value of different sentient creatures, and the relative weights of death (or a statistical chance of death) versus the pain and suffering of laboratory experimentation. Constraints that could be used for this situation include a prohibition on experiments that take sentient life, but that have no potential for life-saving benefits. A second sample constraint could be a prohibition on experiments that cause extreme levels of pain and suffering to sentient beings regardless of the benefits. After constraints are agreed upon, the expected costs and benefits of the various options (including the option of no animal research being done) can be calculated and the utility-maximising option chosen.

Of course, both the constraints and the weightings will be controversial. But even if the exact weightings or utility values would be subject to criticism and debate, the fact is that real decisions involving human interests versus animal interests are made every day. In the case of animal experimentation, these decisions are generally made in favour of human interests at the expense of animal interests because there has been no reconciliation of opposing interests.

The other example used earlier in which a road would be put through a wilderness area would be addressed in a manner very similar to the other two issues. Appropriate constraints would be developed, a social judgment or possibly another approach would be used to arrive at weightings of utility outcomes and weightings of animal/human parties. The probabilities of various outcomes would be estimated based on the best available scientific data. Once probabilities, utilities and weightings are estimated, calculating the value of each option is a simple mathematical exercise and the utility maximising outcome would be recommended. It should be noted that the utility impact of each option to various groups (particularly humans) should be taken on a net basis. In other words, if the benefits of a road in terms of improved access and convenience can be estimated in dollar terms at $10,000,000 and the cost it would take to build the road is $9,000,000, then the net benefit to people would be the $1,000,000 difference, not the full $10,000,000. In many cases the human benefits of development projects only slightly outweigh the human costs, therefore adding animal interests can easily sway the balance in favour of not developing.

VII. CONCLUSION

The difficulties of a pure utilitarian framework are well known. It has been argued here that purely rights-based approaches can lead to just as many difficulties when used to analyse real situations involving difficult trade-offs. An attempt has been made to reconcile the two approaches using constrained optimisation. Many may not be comfortable with this approach, particularly when analysing real dilemmas involving trade-offs. However, this discomfort may stem more from difficulty with the trade-offs discussed than with the approach taken.

It is important to emphasise that these decisions between human interests and animal or environmental interests are made every day. Often these decisions are made in the absence of any system for valuing non-human interests (except at best for the economic value humans place on the environment). And without any framework for weighing their interests, animals and the environment are usually the ones to suffer in the process.

NOTES

[1] One answer that comes to mind is the obligation to spay or neuter an animal. But this obligation has some serious problems that are too complicated to discuss fully here. Briefly, one obvious problem is that if everyone followed this universal obligation, it would result in the end of the population. A second is that it is not an obligation that naturally follows from any more universal principle.

[2] Of course defining 'a risk of death' as a basic interest ignores the difficulties with probabilities previously discussed. But this is cited only because it is the most dramatic contrast in interests. There certainly are other examples where peripheral interests are traded at the expense of more serious ones.

[3] To make this criticism more tangible, let us say that a certain segment of the population enjoys watching animals being tortured. If enough people get even slight pleasure from a TV show featuring this action, then it would outweigh the suffering and possible death of the animal and therefore be considered a morally justified act. The same would hold true if some people enjoyed watching people being tortured.

[4] I would argue that some weighting must be used for the system to reach acceptable conclusions. First, if no weighting is used, then there is no basis for choosing to save a human's life over the life of a mouse or other sentient animal. But to take this one step further, if we were socially engineering the world based on maximising utility, and if the weight of a mouse's happiness and a human's happiness are equal, then it would be better in our utility function to have very few people and very many mice (or other small, sentient animals). This is because it takes much fewer resources to satisfy a mouse and the planet could therefore support many more mice at the expense of one less human. Perhaps we should keep just enough humans around to act as caretakers and maximise food production for the mice. My tendency is to base weighting on a 'mental capacity' criteria

that should only apply between species. A human who is judged to be twice as intelligent (if we could even make such judgments) as another human should not be assumed to be any more sentient. Unlike a weighting which is directly based on mental capacity, using this criteria only between species is not inconsistent in this case because sentience is the real issue, not mental capacity. Mental capacity is only used as a proxy when there is thought to be a correlation with the level of sentience. I propose that sentience, which is not directly measurable, is correlated with mental capacity only for between species differences. For those who would argue that according to the way they interpret sentience, there is no such thing as a 'level of sentience', I say that my definition of level of sentience corresponds to the intensity of the internal experience of suffering, happiness, joy, etc. felt by the individual. I presume that even if a frog for example can feel pain, the richness and intensity of both its positive and negative experience is lower than a human's. However, this is certainly not a presumption that all would readily accept. On the other hand if an argument for marginal cases (for example Pluhar and Rollin 1995 make a strong argument for marginal cases and Dombrowski, 1997 gives a detailed account of the debate over the argument for marginal cases) is accepted, one must reject such a weighting or accept giving a lower weight to mentally deficient humans as well. This is an important debate, but one that need not be resolved here since the system proposed can work with equal or different weightings.

[5] Of course, the decision of just who are the interested parties and how many representatives they should each get in the value judgment process will affect the outcome and could become politically heated in itself. But this does not make such a technique impracticable. Though the appointment process can become highly political, task forces and committees to make value-laden decisions are appointed all the time.

[6] An alternative methodology that can be applied to some situations is to analyse the problem technically using a variety of weighting options. Decision-makers can then be given the break-even points (i.e. the weightings at which one policy choice is functionally identical to another policy choice). The decision-makers then do not need to come up with an exact set of weights; they only need to determine whether the value is higher or lower than the indifference point.

[7] This example is very similar to the actual methodology used in a research project analysing policy options for addressing regional dog overpopulation (Frank, 2001). Thus, the constrained optimisation technique outlined here can be put into practice.

REFERENCES

Anderson, James C. 1993. Species Equality and the Foundations of Moral Theory, *Environmental Values* **2**: 347–365.

Animal Procedures Act, 1986. United Kingdom.

Dombrowski, D.A. 1997. *Babies and Beasts: The Argument from Marginal Cases*. Urbana and Chicago: University of Illinois Press.

Eckersley, Robyn 1998. 'Beyond Human Racism', *Environmental Values* **7**: 165–182.

Frank, Joshua 2001. 'The economics, ethics, and ecology of companion animal overpopulation and a mathematical model for evaluating the effectiveness of policy alternatives', Dissertation, Rensselaer Polytechnic Institute, Troy, New York.

Hammond, Kenneth R., Rohrbaugh, John, Mumpower, Jeryl,and Adelman, Leonard 1977. 'Social Judgment Theory: Applications in Policy Formation', in Kaplan, Martin F. and Schwartz (eds), *Human Judgment and Decision Processes in Applied Settings*. New York: Academic Press.

Lombardi, Louis 1983. 'Inherent Worth, Respect, and Rights', *Environmental Ethics* **5**: 257–270.

Lynch, Tony and Wells, David 1998. 'Non-Anthropocentrism? A Killing Objection', *Environmental Values* **7**: 151–163.

O'Neill, Onora 1997. 'Environmental Values, Anthropocentrism, and Speciesism', *Environmental Values* **6**: 127–142.

Pluhar, Evelyn B. and Rollin, Bernard E. 1995. *Beyond Prejudice: The Moral Significance of Human and Nonhuman Animals.* Durham, North Carolina: Duke University Press.

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

Singer, Peter 1975. *Animal Liberation*. New York: The New York Review.

Taylor, Paul F. 1986. *Respect for Nature: A Theory of Environmental Ethics*. Princeton, NJ: Princeton University Press.

VanDeVeer, Donald 1979. 'Interspecific Justice', *Inquiry* **22**: 55–70.

Wetleson, Jon 1999. 'The Moral Status of Beings who are not Persons: A Casuistic Argument', *Environmental Values* **8**: 287–323.