Logic, Methodology and Philosophy of Science

Proceedings of the Fifteenth International Congress

Hannes Leitgeb
Ilkka Niiniluoto
Päivi Seppälä
Elliott Sober
Editors

22 Biodiversity and triage

MARK COLYVAN *

Abstract. We are currently in the midst of one of the largest mass extinction events in Earth's history and there is no end in sight. Many species have little or no chance of survival and for many others their fate depends on the development of promising, implementable conservation management strategies. These strategies include, for example, captive breeding programs, habitat protection, habitat recreation, and reversing various forms of land degradation. An important step in the development of such conservation strategies is deciding where to best focus conservation efforts. The situation is not unlike an under-resourced war-zone hospital, facing regular massive influxes of casualties. Sadly, it's not possible to even attempt to save all the species currently classified as threatened. Like the war-zone hospital, triage measures need to be implemented to determine where we should spend our time and resources. Such measures are controversial for a number of reasons, not least of which is that they sometimes recommend allowing a particular threatened species to go extinct. I will give a qualified defence of triage, outlining its theoretical underpinnings and discussing its limitations.

Keywords: decision theory, triage, optimisation, conservation management, extinction, climate change.

*University of Sydney, A14 Main Quadrangle, Sydney, NSW, 2006, Australia.
Email: mark.colyvan@sydney.edu.au.

1 The sixth great extinction

There is no doubt that we are in the midst of one of the largest mass extinction events in the Earth's history. This is the sixth such major extinction event and it is clear that humans are largely responsible for this one. Biodiversity is being lost at an alarming rate.[1] According to leading ecologist and conservation biologist Stuart Pimm:

> Current rates of extinction are about 1,000 times the background rate of extinction. These are higher than previously estimated and likely still underestimated. (Pimm et al., 2014, p. 987.)

To put this in perspective, the International Union for Conservation of Nature (IUCN) claims that:

> 17,291 species out of the 47,677 assessed species are threatened with extinction. The results reveal 21 percent of all known mammals, 30 percent of all known amphibians, 12 percent of all known birds, and 28 percent of reptiles, 37 percent of freshwater fishes, 70 percent of plants, 35 percent of invertebrates assessed so far are under threat. (International Union for Conservation of Nature, 2009.)

In light of all this, many species have little or no chance of survival and for many others their fate depends on the development of promising, implementable conservation management strategies. But such conservation efforts need to be strategic.[2]

This raises a crucial question: what is the appropriate conservation strategy or strategies in the face of such an ecological crisis? Several possibilities spring to mind. We could focus conservation efforts on biodiversity hotspots[3] such as Madagascar and the Indian Ocean Islands. Alternatively we could focus conservation efforts on those species facing the greatest threat, wherever they are. Another option is to focus on rare ecosystems that provide unique services. Or we might instead opt for something less systematic and more ad hoc, such as attempting to save threatened species when they come to our attention. Current conservation prioritisation is something of a mix of all of the above and has no doubt met with some success. But there is an argument that we can do better by adopting a triage strategy. To see how this works, we must first rehearse some basic decision theory.

---

[1] Here I'll focus on species-level biodiversity (i.e. diversity of species). Alternatives would be to focus on biodiversity at lower levels such as sub-species or even the genetic level, or at higher levels such as that of ecosystems. See Maclaurin and Sterelny (2008) for a good discussion of how to understand biodiversity.

[2] Although the existence of the sixth mass extinction is not in doubt, there is disagreement about its magnitude. A great deal of the disagreement hangs on the difficult task of estimating the background extinction rate, with which the current rate is compared (Regan, Lupia, Drinnan, & Burgman, 2001; De Voz, Joppa, Gittleman, Stephens, & Pimm, 2015).

[3] These are areas where there is a great deal of biodiversity and which face current or future threats.

2 Conservation prioritisation

2.1 Decision theory

Decision theory is the formal theory of rational choice. It assumes that an agent has a number of actions at her disposal, A_1–A_n, and that the world might be in any number of different states, S_1–S_m (where S_1–S_m form a partition of the ways the world might be, or turn out to be). Outcomes are the result of performing a particular act when the world is in some state or other. That is, an outcome is simply an act-state pair: O_{ij} is the result of the agent choosing action A_i while the world is in state S_j. Probabilities, p_{ij} and utilities u_{ij} are then assigned to each outcome O_{ij}. The expected utility of act A_i is $EU(A_i) = \sum_{k=1}^{m} u_{ik} p_{ik}$, where $\sum_{k=1}^{m} p_{ik} = 1$ We calculate the expected utility of each act at our disposal and the decision rule is: *choose the action with the greatest expected utility, if there is such and action* (Jeffrey, 1990).

A special class of decision problems are known as *optimisation problems*. An optimisation is a decision under constraint where one is attempting to maximise one quantity, subject to a constraint of some kind. For example, consider the problem of enclosing the maximum rectangular area of land, subject to the constraint of a fixed length of fencing. Or, closer to our current purposes, consider the problem of trying to maximise some quantity with a fixed budget. In such situations some outcomes are inaccessible because the costs of the actions involved are too high. Although it is common for the constraint to be financial, it need not be; the constraint can be in terms of any resource. In its most general form, an optimisation problem is one of choosing the action with the greatest expected utility, subject to the constraint in question. Now we apply such methods to our conservation strategies.

2.2 Triage

In war-time hospitals the sick and wounded can arrive in such numbers that not all of them can be given the most thorough treatment. Moreover, some patients are in greater danger of dying than others. Medical staff must make decisions about who gets treated first and, in some cases, who gets treated at all. In the most difficult cases, patients are so seriously injured that little can be done to increase their chances of survival. In other cases, patients could be treated but their chance of survival would not be significantly increased and the treatment would seriously drain resources—both medical supplies and the time of the medical team. Since such resources are limited, decisions need to be made about where best to spend resources. We thus have an optimisation problem of maximising expected lives saved, subject to resource constraints.

The relevance of this to the current environmental situation is obvious enough. We are faced with a situation of limited resources (financial, temporal, etc.) so we cannot do all we would like to preserve all the biodiversity we currently find on our planet.

Suppose we have a ranking of the various threatened species from most to least threatened. We must answer the question: where should we spend our valuable resources? A tempting answer is to focus resources on the most threatened species—those higher up the list. But this is akin to treating the most seriously injured first in our war-time hospital, without considering the chances of success and without looking at the relevant opportunity costs (the lives that otherwise might have been saved). This is a bad strategy in the war-time hospital and it is a bad strategy for saving species.

The answer to our environmental prioritisation problem is to set up and solve the relevant optimisation problem. That is, we need to engage in *environmental triage* (Wilson, McBride, Bode, & Possingham, 2006). Here we try to maximise some appropriate quantity (e.g. reducing the expected number of threatened species), subject to existing resource and budgetary constraints.

It is worth pausing to note that one of the benefits of the decision-theoretic way of looking at such issues is that it forces decision makers to be explicit about what the quantity to be maximised is and how it is to be measured. For example, is the aim to reduce the number of extinctions in the wild, critically endangered species, endangered species or threatened species?[4] The quantity to be maximised (or in this case, minimised) needs to be decided and explicitly incorporated into the set up.[5] We are thus not able to get by with vague and imprecise aims such as "improving the natural environment" or "making the world a better place to live". We may still achieve such lofty goals but we do so by maximising a well-defined, measurable and explicitly-stated quantity such as numbers of species saved from extinction. This approach also forces us to think about the nature of the resource constraints we face. More on this later.

As in the war-time hospital sometimes the result will be that utility is maximised by directing resources to those other than the most serious cases. Often, although there are things that could be done to save a particular species, the optimisation strategy will advise doing nothing for that species. The strategy will advise, in effect, to give up on some species in order to focus resources elsewhere. As a notable defender of triage, ecologist and conservation biologist Hugh Possingham, puts it:

> [w]e should pick winners rather than struggling away with the ones on their last legs. (Possingham, from an interview in the *Sydney Morning Herald*, 6/9/14.)

Just as it must be hard in war-time hospitals to leave some critically ill patients to die, giving up on some of our most endangered species is no easy matter. Leaving species to go extinct is certainly not something we do lightly; it is forced upon us

[4] As is standard, I'm here appealing to the IUCN Red-list classification (International Union for Conservation of Nature 2008).

[5] Technically, it is aways expected utility we are maximising but in practice we often allow more easily measurable surrogates for utility such as money, number of species saved and so forth.

by the overwhelming number of threatened species and the relatively poor prognosis for many of them. Desperate times call for desperate measures. Although the triage model is firmly supported by standard decision theory, it should come as no surprise that triage has attracted some criticism. I will now look at some of this criticism and offer a (qualified) defence of environmental triage against the criticism.

2.3 Limitations of the triage model

There are a number of poorly motivated arguments against triage (e.g. the value of the environment cannot be represented in economic terms or that the environment is sacred and conservation biologists should not play God) but to my way of thinking the best arguments against triage attack the theoretical framework used: optimisation under constraint. In any case, these are the arguments I'll focus on here.

First, we might question the assumption of the triage/optimisation model that the resources are an absolute constraint. It is not always the case that resources are fixed. In at least some circumstances, if the budget or resources are negotiable, it might be better to spend some effort petitioning for a larger conservation budget rather than trying to get by with an unreasonable budget or inadequate resources. There is no room for such negotiations in the standard optimisation models under consideration here.

Another assumption of standard triage models is that there are no costs associated with the reallocation of resources. That is, it is assumed that were we to decide not to pursue conservation efforts in relation to some species, the Orange-bellied parrot (*Neophema chrysogaster*, currently classified as critically endangered), say, we could redirect those resources to conservation efforts for some other species, the Tasmanian Devil (*Sarcophilus harrisii*, currently classified as endangered), say. There are a number of difficulties with this assumption. First, sometimes the funds made available have strings attached. Resources from a bird conservation source may not be allowed to be reallocated to conservation efforts directed at a non-bird. In any case, the conservation efforts themselves are also going to be very different, so a group working on saving a bird may not be qualified to engage in conservation of, say, Tasmanian Devils. At the very least there will be transaction costs. There are also issues of geopolitical strings attached to resources: sometimes funds cannot cross national or state boundaries.

Finally, triage requires a standard probabilistic representation of uncertainty. This is not always available and sometimes probability may not even be the appropriate tool for representing the uncertainty in question (Colyvan, 2008a). For example, standard probability theory, is based in classical logic, which is appropriate only if the world is not gappy. That is, for any proposition *P*, either *P* or not-*P* is the case, even though we may be uncertain about which. But now consider cases where vagueness is involved. Take a question such as 'is a given level of risk of extinction of a particular species acceptable?'. Here the vagueness of the notion of 'acceptable risk' (which permeates much of risk analysis), means that there may not be a fact of the matter.[6] If this is the case, it can be argued that probability theory is not the appropriate tool for the representation of the resulting uncertainties (Regan, Colyvan, & Burgman, 2002).

I take all of these objections to the triage model seriously. I do not, however, take them to be decisive. Rather, they highlight limitations on the applicability of the model. As with any model, it has its domain of useful applications and one should exercise caution when using a model beyond its domain of intended applications. All the above objections draw attention to various idealisations in the triage model that are unrealistic in some circumstances. We can thus respond by acknowledging the idealisations in question and this, in turn, forces us to think about the appropriateness of each of the idealisations in question for the particular application of the model.

For instance, if the budget in question is not fixed, then we simply do not have an optimisation problem. If the budget is genuinely open for negotiation then we should move to another more-appropriate model (e.g. a Nash bargaining model) for the budget negotiation. But once the budget has been negotiated, we can then return to the optimisation model with the newly-bargained budget. Similarly, reallocation expenses can be included in the model—we do not need to assume zero-transaction costs. Dropping this idealisation makes for a slightly more complex model but there is no problem in principle incorporating such complications. Similarly, non-probabilistic representations of uncertainty can be accommodated in decision theory (e.g. vague probabilities Walley, 1991), but again it complicates things. This is not to give up on the basic triage model though. If there is no need to introduce such complications (as will sometimes be the case), then we should stick with the basic model.

So, while these criticisms of the assumptions of the triage model are correct, they do not undermine the basic approach. Rather, focussing attention on the assumptions is useful in determining the appropriate applications of the model and can even help in suggesting modifications of the model and different models for applications beyond the scope of the basic triage model. Now we turn to some more radical attacks on triage.

3 Infinite value

Some have suggested that the natural environment is infinitely valuable. If this were right, saving the natural environment would take priority over mere economic concerns. What is not often noted is that it would also mean any environmental trade-

[6] This, of course, depends on your preferred theory of vagueness. Epistemic accounts of vagueness (Williamson, 1994; Sorenson, 2001) hold that there is an unknowable fact of the matter, but others allow there to be truth-value gaps associated with vagueness (e.g. Fine, 1975). Such vague categories can be sharpened up or, rather, replaced with sharp categories but this brings its own problems (Regan, Colyvan, & Burgman, 2000).

offs would be impossible. Triage would not work but neither would very much else. Thankfully there is no reason to entertain infinite environmental values. The natural environment is valuable but not infinitely so. But lets's look at these issues in more detail.

Douglas McCauley is one who appeals to infinite values in order to prioritise the environment over other concerns.

> I suggest that the aggregate value of a chunk of nature—its aesthetic beauty, cultural importance and evolutionary significance—is infinite, and thus defies incorporation into any ecosystem service programme that aims to save nature by approximating its monetary value. (McCauley, 2006, p. 28.)

The idea is clear enough: if a chunk of nature is infinitely valuable, its value trumps any finite, economic concerns. But what reason do we have for believing that nature is infinitely valuable? McCauley seems to suggest (without further argument) that the infinite value of nature follows from accepting that nature has intrinsic value (McCauley, 2006, p. 28). There are several problems with this line of argument. First, it is anything but clear that nature is intrinsically valuable. And even if it were, it would not follow that its value would be infinite (Justus, Colyvan, Regan, & Maguire, 2009). So what reason is there to entertain infinite value for nature? Of course we can simply stipulate that this is so but this solves nothing. Apart from from being unmotivated, it leaves the door open for others to stipulate alternative uses of land to be infinite: coal-seam gas mines, car parks, and so on. We end up with value inflation by stipulation. If nature is to be seen as infinitely valuable, this needs to be motivated. Without such motivation, this line of argument is without merit. But even setting aside the motivational problem, there are other problems for this view.

The first big problem is that infinite value insufficiently discriminates the salient outcomes (Colyvan, Justus, & Regan, 2010). For example, if we hold that some piece of the natural environment is infinitely-valuable (e.g. mangrove forests), assigning meaningful values to larger regions of that habitat is problematic because there is no extra value to be had in larger areas of it. The problem is that all infinite values (at least in standard decision theory) are equal (and it is not clear that appeals to trans-finite values help here). This approach thus prohibits prioritisation of conservation goals—they are all equal!

The second big problem is that infinite value swamps probability. More formally: we violate the Archimedian condition.[7] Without some clever use of non-standard analysis, non-Archimedian decision theory, or the like, we end up with the expected

[7] The Archimedian condition states that whenever a rational agent prefers A to B and B to C, there exists a real–valued $0 < w < 1$, such that the agent is indifferent between B and the following lottery: a w chance at C and a $(1-w)$ chance at A. Informally: there exists a weighted average of the best and worst outcome which is equal in expected value to the middle outcome. When A is infinite, there is no such w.

utility of any action with a non-zero chance of the desired outcome being infinite. For example, if persistence of an endangered species is considered infinitely valuable, any action with a non-zero chance of the species' survival will have infinite expected value. We would need to be indifferent between hunting the species in question, initiating captive breeding programs, designing reserves for them, poisoning them, and so on, for all of these management strategies have *some* non-zero chance of outcomes that result in the preservation the species in question.[8] But this is clearly the wrong answer; we want to invest in conservation strategies with *higher* chances of success. For example, we want to be able to recommend designing appropriate reserves for rhinos rather than hunting them. Infinite values stand in the way of such common-sense judgements.[9]

It is worth noting explicitly that this does not mean that nature has only finite value. Rather, the lack of discriminatory power we have just seen gives us a pragmatic reason for rejecting appeals to infinite values. But there are also good reasons to reject that the environment has infinite value. Apart from anything else, one needs evidence that something has infinite value; it is not enough to simply assert this as McCauley does. What evidence is there that the natural environment has infinite values? In fact all evidence seems to suggest that its value is finite. If the values in question are instrumental, then we reveal these values via our behaviour. But our behaviour includes making precisely the kinds of trade offs that we have already seen would be impossible were the values in question infinite. The instrumental value of natures is surely finite. One might instead appeal to intrinsic values and insist that the *intrinsic* value of nature is infinite. On this account, value is disconnected from human behaviour so the fact that we behave as though nature is only finitely valuable is neither here nor there. But therein lies the problem with the intrinsic value account: if human behaviour does not reveal the values in question, what does? Intrinsic-value accounts of the value of nature face serious epistemic problems: the values in question are inaccessible and thus there is no evidence that they are infinite.[10] It seems that we cannot coherently let go of our anthropocentric perspective (Grey, 1993).

It might, instead, be claimed that the values in question are incommensurable or that we are fundamentally ignorant of the value of nature. But as Elliott Sober correctly points out, ignorance does not help motivating conservation efforts: "[i]f you are completely ignorant of values, then you are incapable of making a rational decision, either for or against preserving some species" (Sober, 1986, p. 175). If you don't know what something's value is, a fortiori you have no reason to think its value is greater than something else. Similar problems exist for incommensurable value. If something's

[8] This line of objection is due to (Hájek, 2003) where it is raised in relation to the infinite values found in Pascal's Wager.

[9] Some of these problems can be avoided by moving to a non-standard decision theory that allows for comparisons between various infinite outcomes (e.g. Bartha and DesRoches, 2017; Colyvan, 2008b).

[10] Such epistemic problems exist for intrinsic values more generally, either infinite or finite (Justus et al., 2009).

value is incommensurable with, say, economic values, then no trade offs are possible and there can be no rational motivation for allocating budgets to conservation efforts. (Justus et al., 2009.) In short, infinite values, incommensurable values, and unknown values are no basis for conservation.

4 Special pleading

Another line of attack on the triage model is to argue that sometimes, at least, special pleading is in order. Defenders of this line of thought then point to successful cases in support. One such case is sometimes thought to be the California Condor (*Gymnogyps californianus*).

In 1987 the California Condor was on the brink of extinction. At that stage there were only 6 birds remaining in the wild. The bird's cause was relentlessly championed by two conservationists, Noel and Helen Snyder, when others had given up on saving the California Condor (Snyder & Snyder, 2000). The remaining wild Condors were captured and placed in captive breeding programs at San Diego and Los Angeles zoos along with 16 other captive birds. The breeding program was very successful and by 1991 Californian Condors were released back into the wild. By October 2014 there were around 425 California Condors (about 219 in the wild, the rest still in captivity) (U.S. Fish & Wildlife Service, 2014).

Stuart Pimm, for one, presents the Condor case as flying in the face of triage.

> [N]ot all species are equal in their ability to inspire us—the condor is among the most spectacular birds—nor in their ability to extend our applied scientific skills to the limits necessary to save them. Were I endangered, I'd want the Snyders and their colleagues there to ensure I wasn't written off. (Pimm, 2000.)

Presumably, the thought is that because the resources used to save the condor were rather extensive (and hard to justify in cost–benefit terms), triage would have recommended that we should have given up on the California Condor. The success of the Snyders and the captive breeding program thus is supposed to tell against the triage approach.

But it is not clear that the condor case is the problem for triage that Pimm takes it to be. First, there is nothing in the triage approach that suggests that all species are of equal value. For whatever reason, some species might be assigned greater importance (e.g. iconic species), just as in wartime hospital, a General might be given greater weight than a Private. Given the cultural importance of the California Condor, it could well be argued that it should be assigned extra weight in any triage model. If so, the triage model may well recommend the same course of action as the Snyders' captive breeding program. But in any case, the triage calculations would need to be done. One shouldn't simply assume that triage always recommends against assigning resources to critically endangered species.

The second thing to say about this case is that it is misleading to focus on a single successful story. Decision theory does not—nor should it—guarantee the best outcome in a given case. Triage and decision theory in general is probabilistic at its core. This means that sometimes it recommends an action that does not result in the best outcome. But nothing can guarantee the best outcome in any given case. Just as casinos can suffer the odd big loss (but win in the long run), triage will sometimes recommend that the chance of a species recovering are so low as to not warrant investing resources and yet, against the odds, it recovers. The Californian Condor might be such a case.

Moreover, focussing on cases where decision theory does not result in the best outcome is to ignore the many cases where it does get it right. For instance, had the recovery program for the California Condor failed (and the game is not over yet—the Californian Condor is still critically endangered), what would Pimm say about the many species that went extinct in the meantime. But of course the recovery program for the Californian Condor does look to have been successful but the point remains that the resources spent on this single species was at the expense of others. There are always opportunity costs, even in successful cases.

5 Ethics versus decision theory

Triage adopts an explicitly consequentialist framework: the values in question are attached to outcomes, which in turn are products of actions of the agents and the states of the world. Indeed, this is how standard decision theory (including economic theory) proceeds. But it might be argued that this consequentialist framework is controversial. For example, it might be argued that we have duties to care for non-human species and this duty has nothing to do with consequences. At the very least, we ought to give due respect to more explicitly ethical motivations for conserving biodiversity, or so goes this line of thought.

The first thing to note is that although standard decision theory is consequentialist in the sense that it is outcomes that are the bearers of value, it is not consequentialist in any stronger sense. For example, the values attached to outcomes need not be anything to do with promoting the greatest good for the population at large. Indeed, the so-called utility function of decision theory is simply a mathematical function that satisfies the von Neumann-Morgenstern axioms (von Neumann & Morgenstern, 1944), and these axioms are, in many ways, rather liberal. They rule against intransitivity of the utility function[11] but they do not rule against valuing genocide. The theory fixes

[11] If the utility of outcome A is greater then the utility of B, and the utility of outcome B is greater than the utility of C, then the utility of A must be greater than the utility of C.

the structure of the utility function but the values this function takes are left open. In an important sense, decision theory is *amoral*. A natural thought is to appeal to ethics to place further restrictions on the values the utility function can realise (so, for example, we might want genocide to be highly disvalued). On this way of looking at things, at least, there is no conflict between ethics and the decision-theoretic framework that triage appeals to; ethics complements decision theory.

Moreover, any ethical theory cannot "go it alone" so to speak. Ethical theories, such as deontology simply do not give advice about what to do in an uncertain world (Smith & Jackson, 2006). Ethical theories need a decision theory. Now it might be thought that the consequentialist framework adopted by standard decision theory is not well suited to non-consequentialist ethical theories such as deontology or virtue ethics. But non-consequentialist ethical theories can be combined with standard decision theory (Colyvan, Cox, & Steele, 2010). The basic idea, for deontological decision theory, for example, is that further formal constraints are placed on the utility function so that the utility function assigns high values to outcomes arising from certain classes of actions (the "duties") and low values to outcomes arising from another class of actions (the "prohibitions").[12] The details need not concern us; the important point is that ethical theories on their own do not give advice about anything let alone about how best to preserve biodiversity. For this we need decision theory. On this way of looking at things, again there is no conflict between ethics and the decision-theoretic framework; decision theory complements ethics.

In short, ethical theories require decision theory in order to give any practical advice about conservation and non-consequentialist ethical theories can be accommodated in something like the standard decision theoretic account. A version of triage is thus compatible with non-consequentialist ethical theories. There is nothing essentially consequentialist (in the ethical sense) about triage. It is thus hard to see the sense in which triage might be thought to be in tension with ethical sensibilities.

6 Conclusion

I've provided a limited defence of triage as the core decision tool in conservation management and conservation planning. There are devils in the details of its application though. As with models, wherever they are applied, we need to be aware of, and critically examine, the various idealisations and assumptions. Understanding these idealisations and assumptions can help in ensuring that the model is used only in appropriate settings. For the triage model under discussion this means: not accepting

[12] Still, the framework remains consequentialist in that it is outcomes that are the bearers of value. If this is objectionable to the deontologist or to the virtue theorist, they need to develop their own non-consequentialist decision theory. Until that is done, they both remain unable to give practical advice in an uncertain world such as ours.

resource constraints as absolute when they are not; including transaction and reallocation costs when appropriate; employing careful and appropriate representations of the various uncertainties; and employing careful and appropriate representation of utility (which includes appropriate values for iconic species). With these caveats in place, triage is simply the rational way to ensure the optimum allocation of resources for the best conservation outcomes. As I've already stressed, it does not guarantee the best outcome in any single case—nothing can do that—but it does guarantee the best results in the long run. And it is the long run we are interested in.

Finally, a word about complications arising from climate change. One of the most serious threats to the Earth's biodiversity in the medium term is anthropogenic climate change. Various climate mitigation measures, such as carbon sequestration schemes, if managed carefully, may be beneficial for both reducing climate change and for the preservation of biodiversity (Bekessy & Wintle, 2008; Venter et al., 2009). But care needs to be taken here. Too narrow a focus on carbon storage can lead to perverse environmental outcomes. For example, quick-growing plantation forests are a useful measure for reducing atmospheric carbon but they contribute little to biodiversity and may even reduce biodiversity, especially if native grasslands and other native vegetation are cleared for plantation forests. There are also concerns about the effects of plantation forests: the plantation tree may become invasive or otherwise significantly alter ecosystem processes (Putz & Redford, 2009; Lindenmayer et al., 2012). Any plausible strategy for dealing with the preservation of biodiversity must include strategies for curbing the effects of climate change but this must be done in a holistic manner. Environmental goals such as biodiversity and carbon sequestration need to be carefully distinguished and ways of jointly optimising these and other environmental values need to be identified. Triage will be an important part of such holistic strategies but careful attention to the relevant environmental values is crucial to success here.

Acknowledgements. I am indebted to Jack Justus, Hugh Possingham, Helen Regan, and Katie Steele for valuable discussions on the topic of this paper. I am grateful to the audience at the "Future Earth" session at the 15th Congress of Logic, Methodology and Philosophy of Science in Helsinki in August 2015 for some helpful comments and interesting discussion. Work on this paper was supported by an Australian Research Council Future Fellowship grant (grant number: FT110100909).

Bibliography

Bartha, P., & DesRoches, C. T. (2017). The relatively infinite value of the environment. *The Australasian Journal of Philosophy*, *95*(2), 328–353.

Bekessy, S. A., & Wintle, B. A. (2008). Using carbon investment to grow the biodiversity bank. *Conservation Biology*, *22*(3), 510–513.

Colyvan, M. (2008a). Is probability the only coherent approach to uncertainty? *Risk Analysis*, *28*(3), 645–652.

Colyvan, M. (2008b). Relative expectation theory. *The Journal of Philosophy*, *105*(1), 37–44.

Colyvan, M., Cox, D., & Steele, K. (2010). Modelling the moral dimension of decisions. *Noûs*, *44*(3), 503–529.

Colyvan, M., Justus, J., & Regan, H. M. (2010). The natural environment is valuable but not infinitely valuable. *Conservation Letters*, *3*(4), 224–228.

De Voz, J. M., Joppa, L. N., Gittleman, J. L., Stephens, P. R., & Pimm, S. L. (2015). Estimating the normal background rate of species extinction. *Conservation Biology*, *29*(2), 452–462.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, *30*(3–4), 265–300.

Grey, W. (1993). Anthropocentrism and deep ecology. *Australasian Journal of Philosophy*, *71*(4), 463–475.

Hájek, A. (2003). Waging war on Pascal's Wager. *Philosophical Review*, *112*(1), 27–56.

International Union for Conservation of Nature. (2008). Wildlife in a changing world: An analysis of the 2008 IUCN Red List of Threatened Species. *IUCN Website*. Retrieved June 6, 2016, from http://www.iucn.org/theme/species/publications/iucn-red-list-publications

International Union for Conservation of Nature. (2009). Distinction crisis continues apace. *IUCN Website*. Retrieved June 6, 2016, from http://www.iucn.org/content/extinction-crisis-continues-apace

Jeffrey, R. (1990). *The logic of decision* (2nd ed.). Chicago: University of Chicago Press.

Justus, J., Colyvan, M., Regan, H. M., & Maguire, L. A. (2009). Buying into conservation: Intrinsic versus instrumental value. *Trends in Ecology and Evolution*, *24*(4), 187–191.

Lindenmayer, D. B., Hulvey, K. B., Hobbs, R. J., Colyvan, M., Felton, A., Possingham, H. P., … Gibbons, P. (2012). Avoiding bio-perversity from carbon sequestration solutions. *Conservation Letters*, *5*(1), 28–36.

Maclaurin, J., & Sterelny, K. (2008). *What is biodiversity?* Chicago: University of Chicago Press.

McCauley, D. J. (2006). Selling out on nature. *Nature*, *443*(7107), 27–28.

Pimm, S. L. (2000). Against triage. *Science*, *289*(5488), 2289.

Pimm, S. L., Jenkins, C. N., Abell, R., T. M, B., Gittleman, J. L., Joppa, L. N., … Sexton, J. O. (2014). The biodiversity of species and their rates of extinction, distribution, and protection. *Science*, *344*(6287), 987–997.

Putz, F. E., & Redford, K. H. (2009). Dangers of carbon-based conservation. *Global Environmental Change*, *19*(4), 400–401.

Regan, H. M., Colyvan, M., & Burgman, M. A. (2000). A proposal for fuzzy IUCN categories and criteria. *Biological Conservation*, *92*(1), 101–108.

Regan, H. M., Colyvan, M., & Burgman, M. A. (2002). A taxonomy and treatment of uncertainty for ecology and conservation biology. *Ecological Applications*, (2), 618–628.

Regan, H. M., Lupia, R., Drinnan, A. N., & Burgman, M. A. (2001). The currency and tempo of extinction. *The American Naturalist*, *157*(1), 1–10.

Smith, M., & Jackson, F. (2006). Absolutist moral theories and uncertainty. *Journal of Philosophy*, *103*(6), 267–283.

Snyder, N., & Snyder, H. (2000). *The Californian condor: A saga of natural history and conservation*. Princeton: Princeton University Press.

Sober, E. (1986). Philosophical problems for environmentalism. In B. Norton (Ed.), *The Preservation of Species: The Value of Biological Diversity*. Princeton: Princeton University Press.

Sorenson, R. (2001). *Vagueness and contradiction*. New York: Oxford University Press.

U.S. Fish & Wildlife Service. (2014). California condor recovery program. *USFWS*, *31 October 2014*. Retrieved June 6, 2012, from https://www.fws.gov/cno/es/calcondor/PDF_files/2014/Condor%20Program%20Monthly%20Status%20Report%202014-10-31.pdf

Venter, O., Laurance, W. F., Iwamura, T., Wilson, K. A., Fuller, R. A., & Possingham, H. P. (2009). Harnessing carbon payments to protect biodiversity. *Science*, *326*(5958).

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. London: Chapman and Hall.

Williamson, T. (1994). *Vagueness*. London: Routledge.

Wilson, K. A., McBride, M., Bode, M., & Possingham, H. P. (2006). Prioritising global conservation efforts. *Nature*, *440*(7083), 337–340.

Author biography. Mark Colyvan is a professor of philosophy at the University of Sydney and a visiting professor at the Munich Center for Mathematical Philosophy at the Ludwig-Maximilians University in Munich. He holds a BSc(Hons) in mathematics and a PhD in philosophy (the latter from the Australian National University). His main research interests are in the philosophy of mathematics, philosophy of logic, decision theory, and philosophy of ecology and conservation biology. He is the author of *The Indispensability of Mathematics* (Oxford University Press, 2001), *Ecological Orbits: How Planets Move and Populations Grow* (Oxford University Press, 2004, with co-author Lev Ginzburg), *An Introduction to the Philosophy of Mathematics* (Cambridge University Press, 2012), and over 80 refereed papers in philosophy journals, biology journals, and edited collections.