

Wild Animal Suffering is Intractable

Nicolas Delon¹ · Duncan Purves²

Accepted: 10 January 2018 / Published online: 12 February 2018
© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract Most people believe that suffering is intrinsically bad. In conjunction with facts about our world and plausible moral principles, this yields a *pro tanto* obligation to reduce suffering. This is the intuitive starting point for the moral argument in favor of interventions to prevent wild animal suffering (WAS). If we accept the moral principle that we ought, *pro tanto*, to reduce the suffering of *all* sentient creatures, and we recognize the prevalence of suffering in the wild, then we seem committed to the existence of such a *pro tanto* obligation. Of course, competing values such as the aesthetic, scientific or moral values of species, biodiversity, naturalness or wildness, might be relevant to the all-things-considered case for or against intervention. Still, many argue that, even if we were to give some weight to such values, no plausible theory could resist the conclusion that WAS is overridingly important. This article is concerned with large-scale interventions to prevent WAS and their tractability and the deep *epistemic* problem they raise. We concede that suffering gives us a reason to prevent it where it occurs, but we argue that the nature of ecosystems leaves us with no reason to predict that interventions would reduce, rather than exacerbate, suffering. We consider two interventions, based on gene editing technology, proposed as holding promise to prevent WAS; raise epistemic concerns about them; discuss their potential moral costs; and conclude by proposing a way forward: to justify interventions to prevent WAS, we need to develop models that predict the effects of interventions on biodiversity, ecosystem functioning, and animals' well-being.

Nicolas Delon and Duncan Purves have contributed equally to the production of this paper.

✉ Nicolas Delon
nd.delon@gmail.com

¹ The University of Chicago Law School, 1111 E 60th St., Chicago, IL 60637, USA

² Department of Philosophy, University of Florida, 313 Griffin-Floyd Hall, Gainesville, FL 32611-8545, USA

Keywords Ecology · Welfare · Animal suffering · Genetic engineering · Ecosystem management

The Problem of Wild Animal Suffering

Most people believe that suffering is intrinsically bad. In conjunction with contingent facts about our world and a number of independently plausible moral principles, this yields an obligation to reduce or eliminate suffering, except when doing so would violate a more stringent moral requirement or would be overly burdensome. Other things being equal, we ought to refrain from causing pain, and we ought to prevent or relieve pain. This is the intuitive starting point for the moral argument in favor of interventions to prevent wild animal suffering (WAS).

The scope and magnitude of suffering in the world is clearest, some have argued, when we consider the lives of wild animals.¹ If we accept the moral principle that we ought, *pro tanto*, to reduce the suffering of *all* sentient creatures, regardless of species membership, and we recognize the prevalence of suffering in the wild (through disease, starvation, parasitism, predation, disproportionate infantile mortality, etc.),² then we seem to be committed to the existence of a *pro tanto* obligation to intervene in the wild to prevent this suffering.

As proponents of WAS reduction/elimination argue, not only utilitarians are committed to these claims. That's because everyone should believe that suffering is intrinsically bad, regardless of its cause, and that a state of affairs involving less suffering is, other things being equal, better than a state of affairs involving more.³ Even philosophers who do not believe that an instance of suffering is intrinsically bad, where the intrinsic badness is understood as inhering in the state of affairs itself, will nonetheless agree that we have *pro tanto* reasons to prevent suffering. And note that, even though many arguments in favor of reducing wild animal suffering are consequentialist, some non-consequentialist views can be equally demanding regarding the prevention of suffering.⁴

¹ Dawkins (1995), Horta (2010) and Tomasik (2015). Also see an overview at <http://www.animal-ethics.org/population-dynamics-animal-suffering/> (accessed January 18, 2017).

² The prevalence claim might be subject to disagreement or epistemic uncertainty, regarding the scope as well as the intensity of negative experience across species. We set the question aside for the argument's sake.

³ For instance, Nussbaum (2006) advocates an "intelligent, respectful paternalism" (p. 370) and "the gradual supplanting of the natural by the just" (p. 400), and argues that, given our pervasive interference with natural processes and the cruelty inherent in nature, we have positive duties to ensure the flourishing of all creatures with "capabilities"; Sapontzis (1987, Ch. 13) argues that our commonsense principles commit us to accept that, in some circumstances, we have positive duties to rescue "the rabbit from the fox"; Johannsen (2017) (see below) also argues from non-utilitarian premises.

⁴ For instance, Liz Ashford (2003) has argued that Scanlonian contractualism can be very demanding and commit well-off agents to stringent obligations to make significant sacrifices for the global poor. However, Hills (2010) has argued that utilitarianism is uniquely demanding when it comes to WAS, a problem that contractualism has an easier job avoiding.

Of course, competing values such as the aesthetic, scientific or moral values of species, biodiversity, naturalness or wildness, might be relevant to the all-things-considered case for or against intervention. Still, most proponents of interventions to reduce WAS argue that, even if we were to give some weight to such values, no plausible theory could—without undue partiality, inconsistency or implausibility—resist the conclusion that WAS is overridingly important (given its disvalue and prevalence). We are primarily concerned with what we call, interchangeably, massive, systematic, or large-scale interventions. We are not concerned with one-off and/or localized, targeted interventions that do not significantly address the problem as we've presented it. Rescuing a fawn from an icy lake or feeding birds in the winter prevents only a tiny fraction of the total suffering in the wild.

In this paper, we set aside responses to the existence of WAS that acknowledge that WAS is morally bad, but which deny a moral commitment to systematically intervening to prevent it. In making these arguments, authors have appealed, *inter alia*, to rights and species-specific flourishing (Regan 2004; Everett 2001), to the distinction between negative and positive duties in relation to dependence and vulnerability (Palmer 2010), to demandingness (Hills 2010), and to sovereignty and group membership rights (Donaldson and Kymlicka 2011).

Unlike existing responses to WAS, our focus is the deep *epistemic* problem with interventions to prevent WAS. Authors writing about interventions to prevent WAS uniformly concede that there is an important practical constraint on any proposal: an intervention is morally justified only if we are sufficiently confident that the intervention will not make matters worse (Singer 1975; Cowen 2003; McMahan 2010; Sözmen 2013; Horta 2013, 2015; Tomasik 2015; Simmons 2009; Ladwig 2015). In other words, all authors endorse something like the following conditional claim: “‘If we could intervene in the [wild to prevent suffering] without risking ecological disaster, then we should.’” (Johannsen 2017: 4).⁵ We also accept this conditional claim. The crucial question, in our view, is whether the practical constraint is met. We argue that it is not. The indeterministic nature of ecosystems leaves us, at present and for the foreseeable future, with no reason to believe that large scale interventions in the wild would reduce, rather than exacerbate, suffering.

The extent of the epistemic difficulty associated with proposals to prevent WAS has more often been set aside than it has been explored. We begin by considering two large scale technological interventions that have been proposed as holding promise to prevent WAS. We then raise a range of epistemic concerns about these interventions in the following sections. Our concerns derive from an appreciation of the complexity of ecosystem interactions, the effects these interactions might have on the resilience and integrity of ecosystems, and the largely unexplored relationship between ecosystem integrity and wild animal well-being. Because interventions to prevent WAS adversely affect ecosystem resilience and integrity, and because ecosystem resilience and integrity can bear directly on the well-being of animals, we are not justified in believing that interventions to prevent WAS will prevent rather than exacerbate suffering in the wild. We then discuss possible moral costs of interventions to prevent WAS. First, interventions threaten values attached

⁵ Text in parentheses added.

to species, biodiversity, naturalness and wildness. If we take these competing values seriously, we therefore have weighty reasons not to intervene to reduce WAS. Second, interventions to prevent WAS might violate the moral constraint against doing harm. The person, governmental, or nongovernmental entity who deploys harmful interventions to prevent WAS does harm. WAS is something that would merely be allowed by that same agent. If doing harm is morally worse than allowing harm, this is a further powerful moral reason not to intervene until we have a great deal of confidence that intervention will lead to a great reduction in net suffering. The following section addresses some objections to our argument. We conclude by proposing a way forward: to satisfy the practical constraint on interventions to prevent WAS, we need to develop models that reliably predict the effects of interventions on biodiversity, ecosystem functioning, as well as animals' well-being. This leads to a general discussion about how to allocate resources to improve the tractability of important moral causes.

Proposals for Interventions to Prevent WAS

Our critique focuses on two specific proposals for interventions to prevent WAS: gene editing aimed at changing the reproductive behavior of prey (Johannsen 2017, 8) and “reprogramming” predator species to be herbivores (Pearce 2009, 2016). We have chosen these interventions because (a) the proponents of these interventions claim that they do not require continuous widespread human interference in the natural world; and (b) they are representative of the type of intervention that any large-scale attempt to reduce WAS would entail. Demonstrating the epistemic difficulties associated with these specific interventions therefore promises to be illuminating of the general epistemic problem facing interventions to prevent WAS.

Johannsen’s Proposal: Changing the Reproductive Behavior of *r*-Strategists

Johannsen (2017, 8–12) has recently addressed WAS caused by the reproductive behavior of “*r*-strategists”, species of animals that ensure survival of subsequent generations by producing a large quantity of offspring, many of whom die before they reach sexual maturity. The *r*-strategy is used by many smaller vertebrates (and invertebrates) to ensure “survival through quantity”. Offspring do not receive the same careful attention characteristic of larger mammal species who reproduce using the “*K*-strategy”.⁶ As a result, many offspring die prematurely due to predation, disease, exposure, or starvation. Because many of these individuals are sentient, in Johannsen’s words, “The amount of suffering and premature death associated with the *r*-strategy is enormous” (Johannsen 2017, 5).

As an example of the suffering associated with the *r*-strategy, Johannsen describes a population study of the wall lizard (Vitt and Caldwell 2009, 139;

⁶ The *r*/*K*-selection paradigm was introduced by MacArthur and Wilson (1967) and gained traction in the 1970s. Its predictive power and empirical adequacy have since been subject to controversy (see e.g., Reznick et al. 2002). We do not take a stance on this. We treat the distinction as a rough heuristic meant to capture the evolution of different life histories and reproduction strategies.

Barbault and Mou 1988) in which, of 194 reptiles that survived birth for one year, only 48 reached sexual maturity. Johannsen conjectures that the vast majority of the lizards who faced their demise before the age of two suffered “abysmal” fates. We agree this may have been the case for many of them.

How, then, should someone concerned with WAS respond? Johannsen argues that the emerging gene editing technology commonly called CRISPR/Cas9, holds promise to alter wild animals’ reproductive behavior to prevent the massive and painful loss of life suffered by animals like the common wall lizard.⁷ CRISPR can be used to spread nearly any type of genome alteration through sexually reproducing populations. Johannsen hopes that we can use so-called “gene drives”, a CRISPR-based technique to promote the inheritance of a particular gene to increase its prevalence in a population—here, to drive traits that reduce reproductive viability through *r*-strategist populations to reduce their fertility rates. If successful this would reduce the number of offspring that suffer “abysmal” fates, thereby reducing WAS.

McMahan’s and Pearce’s Proposals: Reprogramming Predators

A complement to Johannsen’s proposal has been suggested by McMahan (2010, 2015) and developed in greater detail by Pearce (2009), a prominent advocate of “utopian technologies” to abolish suffering in the biosphere. Pearce suggests that we could genetically “reprogram” carnivores, including through CRISPR, to suppress their inherited predatory traits, while preserving all the other traits required for them to flourish. “The hypernurturing behaviour” of eusocial mammals could be harnessed in carnivores to protect members of species they currently predate; their diet could be replaced by a supply of cultured meat, or we could alter their physiology more radically so that they might subsist on a plant-based diet.

In 2014, a group of experts drew attention to the promise of CRISPR but also to some of its limitations. First, it works only through sexual reproduction, so it cannot affect asexually reproducing species. This includes viruses, bacteria, fungi, many plants, and some animals. Second, there is no way to ensure the continued prevalence of the desired trait in the face of natural selection. Third, again because it operates via sexual reproduction, a significant reduction in WAS could not be achieved for many generations by reprogramming predators, since the beneficial trait would need many generations to spread. This could take “decades or centuries for long-lived organisms.” (Esvelt et al. 2014b; also see Esvelt et al. 2014a). But, setting these limitations aside, the technology seems promising.

⁷ ‘CRISPR’ is an acronym for *Clustered Regularly Interspaced Short Palindromic Repeats*, a naturally occurring feature of certain bacterial genetic codes that functions as a sort of immune defense system against invading viruses. CRISPRs, which contain the genetic code of past invading viruses, are used to help identify and destroy those viruses when they return. The CRISPR/Cas9 technology is based on this remarkable feature of bacteria. It acts as a cutting tool—the Cas9 enzyme is guided to a target DNA sequence and then “cuts it”, either shutting off or activating the targeted gene.

The Epistemic Argument Against Interventions to Prevent WAS

We now turn to the central concern of our paper: the epistemic hurdles facing interventions to reduce WAS by changing the reproductive behavior of *r*-strategists and reprogramming predator species. This section argues that such interventions, when scaled up to meet the actual demands of WAS, are epistemically fraught because of the unpredictability of their interactions with *social ecological systems* (SESs).⁸ Because of the epistemic difficulties associated with interventions to prevent WAS, there is at present, and for the foreseeable future, no reason to believe that the practical constraint on interventions is satisfied.

Basic Ecological Concerns with Changing Species Traits: Resilience

We begin with one simple inductive argument against interventions to prevent WAS that was alluded to in our section “[Proposals for Interventions to Prevent WAS](#)”. Humans have done a poor job anticipating and controlling the effects of large-scale interventions into ecological systems in the past. Consider, for example, the development and implementation of industrial agriculture in the 20th century. High intensity, high yield methods of industrial agriculture have caused diminished topsoil, reduced freshwater availability, aquatic dead zones, biodiversity losses, and large amounts of greenhouse gas emissions. Industrial agriculture has produced “enormous externalized costs, both ecological and social.” (Sandler 2012, 126). Infamous cases of (re)introducing or eradicating species are also cases in point.⁹ Even attempts to control invasive species can have unexpected and unpredictable downstream effects on ecosystems.¹⁰

The epistemic obstacles to successfully implementing interventions to prevent WAS are best illustrated using lessons from contemporary ecology. Many ecologists have recently come to endorse “Resilience Thinking” (Walker and Salt 2006). On this approach, ecological systems are more complex and less stable than outdated equilibrium and succession theories assume. Ecosystems have alternative “regimes” (or “dynamic equilibria” or “basins of attraction”) separated by thresholds. A regime is simply “a region in state space in which the system tends to remain” (Walker et al. 2004). Ecosystems continually undergo adaptive cycles from rapid growth and conservation to release and reorganization, and their variables interact at different temporal and spatial scales. Ecosystems, on this picture, are constantly fluctuating entities. Combined with their processes’ “path dependence” this fluctuation makes ecosystems somewhat unpredictable. An ecosystem’s process is path dependent if, at any given time, it admits of multiple possible outcomes and the probability of future outcomes changes as a function of the system’s initial conditions, history of perturbations, and the particular historical configurations of

⁸ For more on the complexity of SES’s, see our “[Climate Change](#)” section below.

⁹ See Gruen (2011, 185–187) for examples of the “havoc” that our interventions often wreak, e.g., adjudicating between in conflicts between native and non-native species; and Jamieson (2008, Ch. 6).

¹⁰ See for instance Lashley et al. (2018) on the complex ecological consequences of a recent experimental massive culling of feral pigs.

species and other functional groups composing ecological communities (Desjardins 2011). It is possible for ecosystems to unpredictably shift from one regime to another, and some of these regimes will be “degraded states”, which will involve highly variable net sums of suffering.

This new approach to ecological modeling and management does away with the framework of equilibrium and stability, replacing it with the notion of *resilience*, which is understood as the amount of internal or external disturbance or change an ecosystem can absorb while remaining within the same regime, and the speed with which it recovers from disturbance. Resilience thinking involves increasing predictability by promoting resilience in order to avoid, insofar as possible, an abrupt shift from one regime to another. It is particularly important, on this view, to avoid shifts to degraded states with less capacity to absorb environmental changes (Desjardins et al. 2015).¹¹

Ecosystem resilience is a function of a number of interacting features of species and habitats, including structural and response diversity, functional redundancy, modularity, and spatial heterogeneity. As an example of how biodiversity affects resilience, small island ecosystems are very sensitive to biological invasions, and monocrops are very sensitive to proliferating pests. Furthermore, the predictability of the response¹² of communities and ecosystems to perturbations, including global environmental change, is a function of the number of individuals and species. Factors such as trophic chains with fewer levels and fewer species at the highest trophic levels, fragmented or homogenized habitats, and decreased species and gene diversity all negatively impact resilience (Cummings et al. 2005), which in turn can impact their predictability, and yet some argue they would reduce WAS (Tomasik 2017).¹³

In the remainder of this section we highlight some ways in which the interventions we are considering would risk decreasing the resilience of ecosystems, and therefore their predictability.¹⁴

¹¹ Management may not always involve promoting resilience, because resilience is not necessarily desirable (Walker et al. 2004). When a system is in an undesirable regime, good management may require *overcoming* resilience to reach a new regime. Of course, the question of what regimes are “desirable” is up for debate. To make our point here we only need it to be true that the regime state of an ecosystem has significant bearing on the well-being of the ecosystem’s members.

¹² That is, the higher the number of individuals and species, the lower the variance in response, eventually approximating the mean response predicted by theoretical models of the system.

¹³ If Ng (see fn. 30 below) and Tomasik are right that suffering is largely prevalent in nature in part because of *r*-strategists, then interventions that accidentally cause e.g., species extinction or habitat loss would tend to decrease total suffering by reducing the number of *r*-strategists. We have argued against this view because (1) we cannot reliably predict that ecological damage will not have other major impacts and (2), as we argue below, additional considerations tip the scale against taking such risks. The threshold of acceptable risk, of course, depends on the relative importance of these different considerations, including one’s confidence in the disvalue of the existence of *r*-strategists. We also note that this strategy relies on the assumption that these are necessarily lives not worth living. We thank Kyle Johannsen for pressing us on this point.

¹⁴ We do not assume that resilience per se is always desirable. Whether it is depends on many natural and social factors. Our claim is that the descriptive concept of resilience captures one way in which altering ecosystems can affect predictability, which is key to overcoming the practical constraint on interventions we started with.

Interaction Effects and Species Resilience

To see the effect interventions might have on resilience, consider Johannsen's proposal to change the reproductive behavior of *r*-strategists to reduce the amount of suffering on the part of their offspring. There is no question that such an intervention would reduce functional redundancy, which requires different animals living in the same ecosystem that perform the same ecosystemic functions.

As Johannsen acknowledges (pp. 9–10), his solution to the problem of *r*-strategists only begins with reducing fertility of prey species. One must then confront the risk of extinction posed to *r*-strategists by the reduction in fertility, and, relatedly, one must then worry about the effect of reducing the prey population on predator species that rely on them for survival. Changing the reproductive behavior of a species can have unintended effects on (1) the modified species, (2) species that interact with the modified species, and (3) the functioning of the ecosystem of which the modified species is a part.

While the adaptive strategy of *r*-selected species is to rely on the *number* of individuals, *K*-strategists rely on the *quality* of individuals, or their capacity to respond to stress and variations. This gives *r*-strategists flexibility at the population level, making them generally more resilient to major disruptions, even if only at the expense of very many individuals. On the other hand, *K*-strategists are generally more flexible at the individual level. As a population, they depend on the survival and resilience of those who can cope with disruptions. Thus, *r*-strategists are more resilient than *K*-strategists to changes in population dynamics and environmental changes. Changing the reproductive behavior of *r*-strategists to resemble that of *K*-strategists, without increasing flexibility at the individual level, thus risks extinction of the modified species in the face of predation and environmental change.¹⁵ The worry, which Johannsen acknowledges, is that too many of the prey species' young will be predated and the species will be driven to extinction. A related worry is that reduced numbers of prey will cause the members of predator species suffering through starvation, while driving *them* to the brink of extinction as well. Johannsen's solution to both problems is to develop plants that predators can eat in lieu of members of prey species. He says, "Though we won't know for sure until the necessary research is conducted, CRISPR's success with editing plant genes gives us reason to believe that it's possible to develop plants that are suitable for carnivores." (Johannsen 2017, 10). This, in essence, is the converse of Pearce's proposal to reprogram predators to eat plants.

But it is not enough to save prey that we provide predators with the *opportunity* to eat a plant-based diet. They must be programmed to be *disposed* to do so. They must also be programmed not to needlessly engage in hunting behavior. Anyone who lives with domestic cats knows that hunger is not a necessary condition for predatory behavior. Moreover, predation is but one threat to the survival of prey

¹⁵ Some *K*-strategists, the Giant Panda for instance, are inflexible at both the population level and the individual level. The Giant Panda is a threatened species of course, and so it is an exception that proves the rule: flexibility either at the individual level or at the population level is crucial for resilience to disruption. A decrease in flexibility at one level without a corresponding increase in flexibility at the other level can therefore be expected to worsen a species' prospects for survival in the face of disruption.

species. Disease, environmental change, and habitat loss are equally threatening, and a greater reproductive capacity increases *r*-strategist species' resilience against these threats as well. It is unclear how providing carnivores with a plant-based diet alleviates the heightened threat posed to former *r*-strategists by disease or habitat loss. In fact, introducing a plant-based diet for carnivores might threaten prey species more than predation by putting former prey species in competition with predator species for the same resources. Introduced plants might compete for resources with native plants, and they may become invasive, reducing the food available to the former prey species. If the plant-based diet, as is likely, would need to be engineered and massively produced, this would also create an important drain on resources essential to human diets and herbivores.

Gene drives themselves raise environmental concerns. They might have unintended side-effects on the target population. Of course, the complementary technology of "reversal drives" could overwrite unwanted changes. But the technology would not guarantee the reversal of the *ecological side-effects* of gene drives (Oye et al. 2014).¹⁶ The invasive potential of even small-scale experimental introductions of gene-drives has been emphasized in recent work on the population dynamics of wild animals (Esvelt and Gemmill 2017; Noble et al. 2017). There are interaction concerns as well. Any unintended side-effect that is even remotely likely to result from any given gene drive will be potentially combined with any unintended side-effect that might result from other gene drives. Thus, the more comprehensive the response to WAS the greater the likelihood for major missteps.

Ecosystem-Level Resilience

Our first example illustrated the way that interventions to prevent WAS might undermine the ability of a species to adjust to environmental threats. But interventions to prevent WAS can reduce resilience of the *whole ecosystem* as well. As mentioned above, a resilient ecosystem requires a high degree of *functional redundancy*, the presence of multiple species serving the same ecological function. Without functional redundancy, there is no buffer against the downstream effects of the depletion or loss of one species. An ecosystem with less functional redundancy is less resilient and more likely to undergo a regime change. But genetically altering the eating behavior of a predator species necessarily reduces functional redundancy, by removing one channel through which the species' trophic function can be realized. Naeem et al's summary of the considerations one must consider when projecting the effects of biodiversity on ecosystem functioning are illuminating:

each species we introduce possesses functional traits which reflect their tolerances and responses to (e.g., drought or salt tolerance) and impacts on (e.g., nitrogen-fixing or sulfur-reducing) environmental factors such as soil

¹⁶ The authors stress the importance of long-term studies to evaluate the effects of gene drive use on genetic diversity as part of an "integrated management of environmental and security risks". For "any population reduced in numbers will have reduced genetic diversity and could be more vulnerable to natural or anthropogenic pressures." If altered communities became more vulnerable to such pressures, our ability to predict their future states would decrease.

moisture, salinity, and nutrient availability ... The species we introduce will be related to one another by their functional traits, ranging from being nearly redundant ... or nearly singular ... [S]pecies will also possess homologous characters that reflect their shared evolutionary history or phylogeny and will be either closely or distantly related ... Collectively, these many factors determine the biodiversity one finds in a community, all of them influencing flows of nutrients into and out of the inorganic pool, the use and return of water, and the flow of energy sequestered by primary producers and lost through respiration. (Naeem et al. 2009, 5).¹⁷

Altering reproductive or predatory traits of a species that occupies a high trophic level (e.g., apex predators) or plays a “keystone” function can lead to radical changes that reverberate through an ecosystem. Some of these effects include increasing its sensitivity to abiotic perturbations and decreasing the diversity of lower trophic levels. Predators also often occupy the role of “umbrella” species, whose presence supports—as well as it depends on—multiple components of their habitat community. Hence their preservation indirectly supports the preservation of species sharing their habitats (Simberloff 1998). Siberian tigers, for instance, are both a keystone and an umbrella species. Ultimately, the concern is that we might unintentionally shift the ecosystem to alternate, potentially undesirable regimes. Insofar as an ecosystem’s regime state partly determines the well-being of the organisms that constitute it, and insofar as a species’ behavioral traits can affect ecosystem functioning in myriad ways (e.g., primary productivity, nutrient cycling, decomposition, susceptibility to invasion, the spread of disease, or the stability of its populations), answering the question of which genetic alterations will reduce, rather than amplify, WAS is a deeply complicated matter from an ecological perspective.

Clearly, the complexity and unpredictability of SES’s threatens to undermine *any* form of ecological management. But we must bear in mind that the sorts of perturbations that the interventions we are considering are likely to introduce would have management start from a blank slate. That’s because we lack empirically supported models of the range of possible outcomes of such interventions which, by their very nature, are radically novel.

Climate Change

Concerns about the effects on resilience of interventions to prevent WAS are compounded by uncertainty surrounding the effects of climate change. The challenge posed by climate change to the efficacy of interventions to prevent WAS is similar to the challenge posed by climate change to conservation biology. Conservation biologists aim to preserve species *in situ*, in their native habitats. They do this primarily through place-based preservation strategies. For example, one might designate as a protected wilderness area the native habitat of an endangered species. Climate change calls into question whether such place-based preservation can be effective. The trouble is that place-based efforts are incapable of addressing

¹⁷ See Naeem et al. (2009, 5) for numerous citations in support of these claims.

the global changes wrought by climate change. For example, coral reef bleaching cannot be halted by protecting them as marine sanctuaries. This is because the causes of coral reef bleaching are ocean acidification and elevated greenhouse gas concentrations, the origins of which are global rather than local (Sandler 2012, 59). Place-based measures to promote biodiversity are therefore unlikely to “preserve the species’ form of life in their evolved ecological context.” (ibid.)

The interventions to prevent WAS we have considered are inherently place-based strategies as well—they aim to affect the target species *in situ*, changing interactions between individuals and their habitats—and so they face the same challenge. It is impossible to predict with any precision what the effects of climate change will be on a given ecosystem, in part because there are several possible and significantly divergent scenarios. These effects will be determined by the extent of surface air temperature rise; what a warmer world with higher GHG concentrations and more acidic oceans will mean for ecosystem integrity; which mitigation and/or adaptation strategies are effectively adopted by human societies; finally, how human societies respond to different climate change outcomes.

But this is practically impossible to predict because of some of the features of social ecological systems. SES’s are interactive ‘complex adaptive systems’ (a.k.a. ‘CAS’) consisting of bio-geo-physical components and social actors and institutions, including environmental actors, public and private actors, and the public at large. Our understanding of the workings of SESs is limited, and so we cannot know the feedback mechanisms and tipping points endemic to the systems.¹⁸ But it is the *social* component of SESs that makes anticipating the effects of climate change so challenging. The future concentration of greenhouse gas emissions depends on how many emissions occur, which in turn depends on future human decision-making and its impact on economic developments, policy-making, and technological innovation. But, and this point is crucial, the way human societies use technologies is impossible to predict *ex ante*. Dale Jamieson’s remarks about the unpredictable nature of technology are aptos:

[W]hen it comes to performing the benefit–cost calculation concerning many technological innovations, we are ignorant rather than uncertain. Not only can we not do the sums, but we don’t know with any precision what sorts of things we would have to know in order to do them... [W]hile scientific information would be important to performing the benefit–cost calculation, it is information about the actual human use and social deployment of the technologies in question that matters most. The benefit–cost calculation for technological innovations depends on who will control them, who has access to them, how they will impact on the global (mal)distribution of wealth, how they will affect human health, how they will affect animal welfare, and on their ecological consequences. ... For any particular technological innovation we are quite ignorant about most of this, in part because the actual benefits and costs of the innovation will depend on decisions that people make subsequent

¹⁸ Construing ecological systems as SESs also entails that promoting resilience, and therefore managing unpredictability, does not require a hands-off approach but rather an integrated form of management about which we say more later.

to the initial decision to deploy the technology. What humans will do in the future is not just unknown ... it is indeterminate.¹⁹ (Jamieson 2002, 312–3)

Given the political, social, and economic upheaval of the last few years, it is safe to say that we have very little idea of what climate scenario to expect in 2100, and even less of an idea of what the human response will be. The one thing that is clear is that, whatever emissions scenario we envision, we can expect more, and more rapid, ecological variability than we have seen in recent history (Sandler 2012, 52). The upshot for WAS is that any technological intervention to prevent WAS must grapple with the possibility that the ecosystems of the species they are targeting will be radically, and unpredictably, different, 20, 50, or 100 years from now.

This exacerbates concerns raised earlier against Johannsen's proposal to prevent WAS by reducing the number of offspring of *r*-selected species. For, as we saw, populations of species whose members have few offspring and longer gestation periods are less likely to be able to adapt to ecosystem changes than populations whose members rapidly produce large numbers of offspring (Sandler 2012, 55). Given relatively stable ecosystem conditions, perhaps we can anticipate and respond to these changes in advance, but climate change calls the feasibility of such responses into question.

Indifference

In addition to the low tractability of addressing WAS itself, advocacy has proven unpersuasive. Recall that part of “Resilience Thinking” consists in seeing ecosystems as enmeshed within SESs. Given the current resistance of the environmental community, and likely most laypeople, to massive interventions, we seem, presently at least, nowhere near the stage where even beginning to experiment would receive enough public support, funding, and scientific resources. This distinguishes advocacy to reduce WAS from advocacy to reform or abolish factory farming, which, in addition to being more tractable, is more likely to be supported by the public. Fewer competing values, except perhaps for convenience and affordability, prevent people from endorsing farm animal welfare advocacy.

This lack of public support is likely due to a number of factors. One factor is that people are insensitive to large numbers. *Scope insensitivity* or *scope neglect* is a cognitive bias such that, as numbers increase exponentially—for instance the number of birds one can rescue from an oil spill—people's willingness to pay only increases linearly. In fact, sometimes people may be willing to give more to save few individuals than to save a large number of individuals. The phenomenon may be irrational, but it is widespread and well known by economists and advocates. It is also described as the result of “psychic numbing” (Dickert et al. 2015; Slovic 2007). The more lives are at risk, the smaller the emotional response. Relatedly, people are susceptible to an ‘identifiable victim effect.’ This is important when we

¹⁹ In a context of high decision stakes and/or high system uncertainty, the relation between policy and science may be helpfully approached from the standpoint of “Post-Normal Science,” in particular given the conflicting values involved (Funtowicz and Ravetz 1990).

are considering the public response to WAS,²⁰ or any other sort of animal suffering for that matter. Brian Tomasik provides the following rough estimates of animals in the wild: 60 billion land birds, 100 billion land mammals, 1 trillion land reptiles, 1 trillion land amphibians, 10 trillion fishes.²¹ Tomasik also believes that the majority of lives in the wild, especially those of *r*-selected species (the overwhelming majority) are not worth living, or at least involve egregious suffering. In aggregate, this amounts to a mind-boggling sum of suffering. These are staggering numbers, but most people are unlikely to be swayed by them, because they have a difficult time comprehending the scale of suffering, owing to the abovementioned biases.

Furthermore, people are generally too ill-informed to understand the implications of emerging technologies. Rightly or not, the social feedback of massive technological interventions in a domain that people tend to conceive either anthropocentrically (e.g., as an economic, scientific, or aesthetic resource) or as a source of intrinsic value is unlikely to tap into people's most acute moral sensitivity. In other words, psychic numbing and people's lack of scientific literacy or distrust for emerging technologies create hurdles for advocacy to reduce WAS that compound the intractability of the initial problem. Insofar as these are potent factors, they render supposedly beneficial interventions in the wild even more intractable. As a result, such interventions will be even more difficult to defend publicly.

Competing Values and the Moral Constraint Against Harming

We have argued that the complexity of ecosystems, the unpredictability of climate change, and the indeterminacy of human behavior together entail that interventions to prevent WAS are as likely to exacerbate the problem of WAS as they are to ameliorate it. In this section we discuss the normative implications of our epistemic conclusion. One normative implication is already clear. With respect to the prevention of suffering, there is presently no reason to favor interventions to prevent WAS. But there are additional reasons for concern that, at present, tip the moral balance *against* intervening.

First, interventions to prevent WAS, successful or not, will cause harm. Whether we reprogram predators or reduce the fertility of *r*-strategists, individuals will be caused to suffer harms (painful death, disease or some other terrible fate) they would not have suffered had we not intervened.²² Whatever institution or individual administers the intervention is therefore responsible for doing harm in a

²⁰ Advocates to reduce WAS are aware of this and now tend to focus their messages on fewer, identifiable victims. See e.g., Animal Ethics: <http://www.animal-ethics.org/>. Accessed March 1st, 2017.

²¹ Tomasik, <http://reducing-suffering.org/how-many-wild-animals-are-there/> Accessed March 1st, 2017.

²² Here Derek Parfit's non-identity problem (1976, 1984) rears its head, as it does with any policy or action that alters procreative behavior. The core of the problem is how to morally assess actions that both affect the quality of life of future individuals and determine which individuals will come to exist in the future. Here the concern is that, as long as future wild animals have lives worth living, and they would not have existed had we not intervened to prevent WAS, then it is not clear that we can identify a sense in which our intervention harms them. Addressing the non-identity problem, or its many purported solutions, is beyond the scope of this paper.

way that they are not responsible for harm caused by their *failing* to intervene. Harm caused by an agent's non-intervention in the wild is harm (e.g., harm caused by predation or premature death) that the agent merely allows. The moral significance of the doing/allowing distinction is hotly contested, and we will not settle that contest here.²³ We merely point out that the doing/allowing distinction is a central feature of "common sense" morality. It is also important for maintaining a space of permissible options for moral agents in order to avoid an overly demanding moral theory. As David Morrow puts the point in a paper describing the implications of the distinction for the ethics of geoengineering, "given the kinds of harms we are morally forbidden from causing, a requirement to prevent similar harms from befalling others would demand more of our time, energy, and resources than can reasonably be expected of us" (2014, 30). The moral significance of the doing/allowing distinction is also supported by common objections to consequentialist moral theories, including the implication that a surgeon might be morally required to kill one healthy patient in order to use their organs to save several other patients in need of transplants.²⁴

If doing harm is morally worse than allowing harm, then the harm that an agent does by intervening in the wild is morally worse than the harm that the agent merely allows by not intervening in the wild.²⁵ On a moderate view of the moral importance of the doing/allowing distinction, this means that interventions to prevent WAS would be morally justified only if they *prevented a great deal* more harm than they caused. There is some number of lives n such that we would be permitted to kill one to save n lives. But, given all we know, this is not the position we are in with respect to interventions to prevent WAS.²⁶

But suppose the doing/allowing distinction has no intrinsic moral significance. Even still, we might have instrumental reasons for adhering to a somewhat similar distinction. Savulescu and Persson (2012) argue we should reject "the act-omission doctrine of common-sense morality" because "as our powers of action increase [due to the inventions of scientific technology], so does the range of what we let happen through failures to use these powers" (p. 60) But they caution against being too eager to use our newfound technological powers to attempt to actively improve the world. This is because they emphasize that "it is generally much easier to harm than to benefit." To take a simple example, there are many more ways to kill many people while driving a car than there are ways to save many lives. There are exceptions to this general rule, but interventions in ecosystems are not one of them. They write:

²³ Key figures in this debate include Rachels (1975), Foot (1967), Thomson (1976) and Bennett (1995). Kagan (1989) provides perhaps the most compelling sustained attack on the doing/allowing distinction.

²⁴ Thomson (1976, 206).

²⁵ David Morrow (2014) defends a similar claim in the context of discussing the ethics of geoengineering the global climate system in response to climate change. He argues that, because of the existence of a moral asymmetry between doing and allowing harm, "the bar for establishing [geoengineering's] potential moral permissibility is higher than some scientists believe." (124). The moral relevance of the distinction between "a negative duty not to harm" and a "positive duty to assist" also forms the backbone of Clare Palmer's defense of the "Laissez-faire Intuition", the intuition that there is no presumptive duty to assist wild animals (2010, 75).

²⁶ This objection does not apply to calls to *refrain* from reintroducing predator species to areas where they have been driven to extinction. Reintroducing such species would count as *doing harm*.

it is much more difficult to improve significantly upon comparatively well-ordered ecological systems on Earth than it is to damage them seriously.... Hence, great caution is in place when we contemplate large-scale interventions with ecosystems, since unforeseen effects are likely [and] are likely to be mostly for the worse. (p. 46)²⁷

Furthermore, we know with some degree of certainty that interventions to prevent WAS will directly threaten non-welfare values that many people take seriously. For example, some environmentalists hold that wild nature has significant (moral or aesthetic) value. Others hold that species' value depends on their historical evolutionary origins. On both of these views, the value of an ecosystem or species is necessarily undermined by human interference.²⁸ Using CRISPR to directly change the genomic composition of organisms undermines their natural value (if anything would). Of course, proposals to reduce WAS are strictly welfarist, and those are controversial values. Yet, whether or not they are objectively important values, the fact that many people take them seriously must be accounted for in a comprehensive assessment of the impact of interventions to prevent WAS.

It is also arguable that suffering plays an (extrinsically) important role in human lives, contributing to their overall desirability. Perhaps suffering even contributes (extrinsically) to our finding nature valuable. Some have argued that a genuine appreciation of nature entails an affirmation of predation and suffering (Hettinger 1994; Rolston 1988). The extrinsic value of suffering and the non-welfare values people care about could thus make policies aiming to eradicate WAS more “transformative” than we expect. Human preferences and behaviors in SESs, and how they interact with the ecological ‘services’ and socioeconomic activities afforded by the system, are hard to anticipate. WAS-reducing policies would affect our values in ways that we cannot anticipate before experiencing the transformation they entail. Hence, we cannot fully anticipate whether these changes would be desirable.²⁹

Finally, interventions to prevent WAS require intellectual and material resources. Advocacy groups fundraising for such interventions would likely divert money away from other causes that may be more tractable (e.g., issues like the abolition or reform of factory farming). This is why the meta-charity Animal Charity Evaluators does not (yet) recommend prioritizing donations toward WAS—mainly because of its low tractability—even though they recognize the importance of the problem.³⁰

In the next section, we briefly address a number of objections to our epistemic argument.

²⁷ It is easier to damage a complex system than to improve it because the ways in which “[it] could be damaged are indefinitely more numerous than the ways in which it could be improved” (Savulescu and Persson 2012, 51).

²⁸ Rolston (1989), Katz (2000) and Cafaro (2001). Sandler (2012) summarizes these different accounts of the value of species.

²⁹ We are appealing, rather loosely, to the account of transformative experience developed by Paul (2014).

³⁰ Note that one of the early proponents of “welfare biology”-based interventions to reduce the suffering of wild species, a recurrent reference in the WAS community (Ng 1995), has recently argued that we should focus on the suffering of domesticated species before dedicating significant resources to WAS (Ng 2016).

Objections

One might object to our arguments thus far as follows: although climate change is expected to exacerbate uncertainty, it is also likely to cause widespread harm to a great deal of sentient animals. Even if, on balance, we don't know that climate change will increase overall suffering, the fact that it is likely to cause widespread harm would seem to encourage intervention where we are confident we can do more good than harm. This objection is toothless, however. We have argued precisely that we cannot be confident that interventions such as reprogramming *r*-strategists and reprogramming predators will reduce, rather than exacerbate, the harms associated with climate change.³¹

We will now consider a handful of objections raised by Cowen (2003) to the sort of epistemic argument we have offered:

- (1) We are already intervening in nature all the time without knowledge of (or concern for) our effects.³² Interventions to prevent WAS don't add to that uncertainty and so should be considered on their own merits. Our argument simply implies uncertainty about all policies, including non-intervention.

In reply, we admit that, against the backdrop of global climate change and social uncertainty, maybe interventions are just noise and hence do not exacerbate unpredictability. But, certain types of interventions *will* generate more unpredictability by decreasing resilience and undermining ecosystem functions that preserve integrity in the face of climate change. Moreover, that we are presently intervening in nature with great ignorance of, and indifference to, the effects of our intervention does not imply that we are justified in doing so to a greater extent.

But our epistemic concern about climate change and uncertainty might seem to apply to all interventions and therefore appear too strong. It is sometimes permissible, sometimes required, to intervene to preserve species, reduce pollution, manage populations, or curb carbon emissions, even though we cannot be fully confident that such actions will not have devastating long-term consequences. So why would the epistemic argument block interventions to reduce WAS but not others? For example, suppose we want to assist the migration of Pika to a northern ecosystem where they can escape the existential threat posed to them by global warming. But we cannot predict the effects of climate change on this new habitat in part because we cannot predict whether a country will unilaterally develop and deploy geo-engineering to combat climate change. If a country deploys geo-engineering, this might alter rainfall patterns in ways that make the new habitat unviable for the Pika. Here uncertainty surrounding the social response to climate change undermines the efficacy of assisted migration. So we are not in a position to know whether assisted migration is better or worse than no migration. The epistemic problems associated with climate change threaten the rationale of *any* large scale ecological intervention.

³¹ We thank Glen Miller for these suggestions.

³² Also see Nussbaum (2006, 366–370).

First, we concede that our argument might undermine the justification of a variety of interventions into the wild, not only those intended to prevent WAS. But non-welfare values might tip the balance in favor of some interventions rather than others. For example, if we intrinsically value ecosystem health or the natural-historical value of species, then some interventions can be justified if they *increase ecosystem resilience*, even if they are not expected to reduce suffering. Because a resilient ecosystem can absorb greater external disruption, ecosystem health is promoted by resilience even in the face of climate change. For this reason, removing non-native predators might be a justified intervention in the wild if it increased resilience (e.g., by increasing functional redundancy). While a concern for the prevention of suffering cannot provide us with guidance in the face of climate change, a concern for other values might.³³

(2) Small-scale interventions are unlikely to lead to ecological disaster. Couldn't we do those?

Yes, we could! Small-scale interventions that we can be confident will not have large-scale unforeseeable consequences are not our concern. Freeing a particular deer caught in a particular fence is unlikely to lead to abrupt and unexpected regime change. The interventions we have critiqued are not small-scale. If they were, they would not significantly address WAS. An interesting and optimistic possibility is that, given the scale of WAS, even *minor* improvements (e.g., genetic or environmental) to wild ecosystems might have dramatically positive consequences. They might, but, assuming their ramifications are wide indeed, the question remains: Are we confident enough that they will do more good than harm?

(3) Eliminating the European wolf was clearly a good thing from the perspective of suffering and did not lead to ecological disaster. There was no impending ecological problem when wolves were reintroduced to Yellowstone.

This is in fact highly contentious. The reintroduction of the Yellowstone wolf arguably had a number of benefits with respect to promoting animal well-being, and while the elimination of the European wolf did not lead to disaster, we have no idea whether it lowered or increased overall suffering. *That* it is highly contentious is further support for our concern that the welfare-based defense of interventions to prevent WAS is tenuous.

In sum, we agree that many small-scale interventions might be beneficial, and that such interventions might gradually make WAS locally less intractable. However, the interventions we have considered are large-scale; they can be expected to have large-scale consequences, for better or worse.

³³ We thank Glen Miller for pressing us on this issue.

Adaptive Management: A Way Forward

The upshot of the discussion thus far is that we need more time and small-scale experiments before we can start actively intervening at gradually growing scales. Indeed, this is how conservationists and restoration ecologists deal with uncertainty. A consequence of Resilience Thinking, Adaptive Management is a learning process whereby hypotheses, models, and policies are continuously adjusted in response to experimental ecological interventions and unpredicted outcomes. It emphasizes, writes Hourdequin (2017, 452), “that our knowledge of ecological systems is partial, incomplete, and evolving, and it integrates responsiveness to new information or changing conditions. ... [E]ffective adaptive management is a recursive process, in which strategies are tested, evaluated, and modified based on new knowledge, then tested and evaluated again.”

Our arguments in this paper are mainly concerned with the present tractability of possible interventions. However, our ignorance with respect to crucial features of SESs might undermine, not just the tractability of the WAS cause, but its ‘meta-tractability’—the probability that additional resources devoted to *researching* the problem could improve the tractability of the cause. In light of current ecological science, we are tempted by the view that the problem is *deeply* intractable in a way that sheds doubt on the marginal utility of further research. This is because WAS is a *global* problem, which no aggregate of local interventions can address without encountering the hurdles we have described. For example, experimenting on an isolated population of rodents living at high elevations on a sky island runs smaller ecological risks. But it doesn’t even begin to address the global problem. If, on the other hand, comparable islands were somehow interconnected, the epistemic problem would resurface.

Still, we leave open the possibility that further research might make WAS more tractable. How then can we gain the information required to engage in responsible interventions, when all we have are abstract, idealized mathematical models? *Robustness analysis* can help here, namely, a formal method to determine which models make trustworthy predictions (Weisberg 2006). Even with very different assumptions, some scientific models—for instance models of prey-predator dynamics—can lead to strikingly similar results, following a common structure and underlying principles (e.g., the Volterra Principle). In such cases, *even before actual empirical observations are made*, we are likely to have a robust theory of the system we aim to describe.

In the face of uncertainty or ignorance, ecological management can benefit from robust theories, and so could interventions to reduce WAS, provided we could incorporate well-being, or some proxy for it, as one metric among others used in the models. Performing robustness analysis on unique SESs is challenging. Levins (1966) has argued that, given the practical constraints to which both observation and computation are subject, the idealized models that scientists use must make tradeoffs among three desirable features: *precision*, *generality*, and *realism*. By reintroducing particulars that make models more realistic, for instance, one loses in generality, potentially creating models too complex to be of any use. On the other

hand, models generating hypotheses about large-scale ecological and evolutionary patterns are too simple to generate precise predictions about particular systems (Barker and Odling-Smee 2014, 204), let alone the sort of precision we would need to monitor WAS. Applied to WAS, a robustness analysis approach suggests that a given model will be sufficiently accurate in predicting the effects of technological interventions to prevent WAS for a *particular* SES, but the same model cannot be generalized to *other* systems without significant tailoring. We believe this approach reflects a nice compromise between setting the epistemic bar too low (thus risking ecological disaster) and setting it too high (thus risking avoidable moral catastrophes). The compromise also avoids the charge of *status quo* bias—our ingrained tendency to leave, and prefer, things as they are.

With these limitations in mind, and given the scope and scale of the alterations envisioned by proponents of interventions to reduce WAS, epistemic, but also practical and regulatory, hurdles loom large. A matter of centuries starkly diminishes the prospect of reliable predictions regarding the overall consequences of such interventions.

Conclusion

We argued that large-scale technological interventions to reduce suffering in the wild face important epistemic hurdles. Ecological science and facts about human agency suggest that some ecological systems are not predictable in the way that would make such interventions morally permissible, especially in the context of global environmental change. We examined two types of genetic engineering, made possible by gene drive technology utilizing CRISPR, to alter, respectively, the reproductive behavior of *r*-selected species and the behavior of predators. We have argued that each intervention might have unforeseen ramifications the value of which cannot be properly assessed. The trouble is that the partial indeterminacy of SESs and the extreme complexity of the ecological processes captured by existing models make the cascading effects of large-scale interventions extremely hard to predict. The problem is compounded by the rate, magnitude and uncertainty associated with global environmental change. We have also argued that interventions to reduce WAS face more serious epistemic challenges than other forms of ecological management and interventions in nature.

A possible compromise, we have suggested, is to take the cue from Adaptive Management and test interventions on a piecemeal basis, gradually adjusting their design and implementation in response to the resulting ecological changes (or lack thereof). We can now break down our view into two theses. Our *negative thesis* is that we are not justified in thinking that existing proposals would in fact reduce total suffering—we are just in no position to know. Our *positive thesis* is that, for interventions to be justified, we need first to develop metrics and models that predict the effects of interventions (or any human activity) on biodiversity, ecosystem functioning, and animals' well-being. Whether it is possible to develop these metrics and models is a question we wish to leave open for now.

Acknowledgements This article benefitted from fruitful early discussions with Dale Jamieson, detailed comments from Kyle Johannsen, Glen Miller, and two anonymous referees, and feedback from audiences at the Bovay Workshop on Engineering and Animal Ethics at Texas A&M and the CRE/GREA Conference on Animal and Environmental Ethics at McGill University in Montreal. We also thank Clare Palmer and Gary Varner for their encouragement and suggestions.

References

- Ashford, E. (2003). The demandingness of Scanlon's contractualism. *Ethics*, 113(2), 273–302.
- Barbault, R., & Mou, Y. P. (1988). Population dynamics of the common wall lizard, *Podarcis muralis*, in southwestern France. *Herpetologica*, 44, 38–47.
- Barker, G., & Odling-Smee, F. J. (2014). Integrating ecology and evolution: Niche construction and ecological engineering. In D. Barker et al. (Eds.), *Entangled life: Organism and environment in the biological and social sciences* (pp. 187–211). Dordrecht: Springer.
- Bennett, J. (1995). *The act itself*. New York: Oxford University Press.
- Cafaro, P. (2001). Thoreau, Leopold, and Carson: Toward an environmental virtue ethics. *Environmental Ethics*, 23(1), 3–17.
- Cowen, T. (2003). Policing nature. *Environmental Ethics*, 25, 169–182.
- Cummings, K. W., Barnes, G., Perz, S., Schminck, M., Sieving, K. E., Southworth, J., et al. (2005). An exploratory framework for the empirical measurement of resilience. *Ecosystems*, 8, 975–987.
- Dawkins, R. (1995). God's utility function. *Scientific American*, 273, 80–85.
- Desjardins, E. (2011). Historicity and experimental evolution. *Biology and Philosophy*, 26, 339–364.
- Desjardins, E., Barker, G., Lindo, Z., Dieleman, C., & Dussault, A. (2015). Promoting resilience. *The Quarterly Review of Biology*, 90(2), 147–165.
- Dickert, S., Västfjäll, D., Kleber, J., & Slovic, P. (2015). Scope insensitivity: The limits of intuitive valuation of human lives in public policy. *Journal of Applied Research in Memory and Cognition*, 4(3), 248–255.
- Donaldson, S., & Kymlicka, W. (2011). *Zoopolis: A political theory of animal rights*. Oxford: Oxford University Press.
- Esvelt, K., Church, G., & Lunshof, J. (2014b). 'Gene Drives' and CRISPR could revolutionize ecosystem management. *Scientific American*, July 17, 2014. Published online. <https://www.scientificamerican.com/author/kevin-esvelt-george-church-and-jeantine-lunshof/>. Accessed December 15, 2017.
- Esvelt, K., & Gemmill, N. J. (2017). Conservation demands safe gene drive. *PLoS Biology*, 15(11), e2003850. <https://doi.org/10.1371/journal.pbio.2003850>.
- Esvelt, K., Smidler, A., Catteruccia, F., & Church, G. (2014a). Emerging technology: Concerning RNA-guided gene drives for the alteration of wild populations. *eLife*, 3, e03401. <https://doi.org/10.7554/eLife.03401>.
- Everett, J. (2001). Environmental ethics, animal welfarism, and the problem of predation: A Bambi lover's respect for nature. *Ethics and the Environment*, 6(1), 42–67.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 5–15.
- Funtowicz, S. O., & Ravetz, J. R. (1990). *Uncertainty and quality in science for policy*. Netherlands: Kluwer Academic Publishers.
- Gruen, L. (2011). *Ethics and animals: An introduction*. New York: Cambridge University Press.
- Hettinger, N. (1994). Valuing predation in Rolston's environmental ethics: Bambi lovers versus tree huggers. *Environmental Ethics*, 16(1), 3–20.
- Hills, A. (2010). Utilitarianism, contractualism, and demandingness. *Philosophical Quarterly*, 60(239), 225–242.
- Horta, O. (2010). Debunking the idyllic view of natural processes: Population dynamics and suffering in the wild. *Telos: Critical Theory of the Contemporary*, 17(1), 73–90.
- Horta, O. (2013). Zoopolis, intervention, and the state of nature. *Law, Ethics, and Philosophy*, 1, 113–125.
- Horta, O. (2015). The problem of evil in nature: Evolutionary bases of the prevalence of disvalue. *Relations*, 3, 17–32.
- Hourdequin, M. (2017). The ethics of ecosystem management. In A. Thompson & S. Gardiner (Eds.), *Oxford handbook of environmental ethics* (pp. 449–462). Oxford: Oxford University Press.

- Jamieson, D. (2002). Discourse and moral responsibility in biotechnical communication. *Morality's progress* (pp. 308–320). Oxford: Oxford University Press.
- Jamieson, D. (2008). *Ethics and the environment: An introduction*. New York: Cambridge University Press.
- Johannsen, K. (2017). Animal rights and the problem of *r*-strategists. *Ethical Theory and Moral Practice*, 20(2), 333–345.
- Kagan, S. (1989). *The limits of morality*. Oxford: Oxford University Press.
- Katz, E. (2000). The big lie. In W. Throop (Ed.), *Environmental restoration* (pp. 83–93). Amherst, NY: Humanity.
- Ladwig, B. (2015). Against wild animal sovereignty: An interest-based critique of Zoopolis. *Journal of Political Philosophy*, 23, 282–301.
- Lashley, M. A., Jordan, H. R., Tomberlin, J. K., & Barton, B. T. (2018). Indirect effects of larvae dispersal following mass mortality events. *Ecology*, 99(2), 491–493. <https://doi.org/10.1002/ecy.2027>.
- Levins, R. (1966). The strategy of model building in population biology. *American Scientist*, 54(4), 421–431.
- MacArthur, R. H., & Wilson, E. O. (1967). *The theory of island biogeography*. Princeton: Princeton University Press.
- McMahan, J. (2010). The meat eaters. *New York Times: The Opinionator*. September 19, 2010.
- McMahan, J. (2015). The moral problem of predation. In A. Chignell, T. Cuneo, & M. Halteman (Eds.), *Philosophy comes to dinner: Arguments about the ethics of eating* (pp. 268–294). London: Routledge.
- Morrow, D. (2014). Starting a flood to stop a fire? Some moral constraints on solar radiation management. *Ethics Policy and Environment*, 17(2), 123–138.
- Naeem, S., Bunker, D., Hector, A., Loreau, M., & Perrings, C. (2009). Introduction: The ecological and social implications of changing biodiversity. An overview of a decade of biodiversity and ecosystem functioning research. In S. Naeem, D. Bunker, A. Hector, M. Loreau, & C. Perrings (Eds.), *Biodiversity, ecosystem functioning, and human wellbeing: An ecological and economic perspective* (pp. 1–13). Oxford: Oxford University Press.
- Ng, Y.-K. (1995). Towards welfare biology: Evolutionary economics of animal consciousness and suffering. *Biology and Philosophy*, 10(3), 255–285.
- Ng, Y.-K. (2016). How welfare biology and commonsense may help to reduce animal suffering. *Animal Sentience*, 7(1). <http://animalstudiesrepository.org/animations/vol1/iss7/1/>. Accessed June 1, 2017.
- Noble, C., Adlam, B., Church, G., Esvelt, K., & Nowak, M. (2017). Current CRISPR gene drive systems are likely to be highly invasive in wild populations. *bioRxiv* 219022. <https://doi.org/10.1101/219022>.
- Nussbaum, M. (2006). *Frontiers of justice: Disability, nationality, species membership*. Cambridge, MA: The Belknap Press.
- Oye, K. A., Esvelt, K., Appleton, E., Catteruccia, F., Church, G., Kuiken, T., et al. (2014). Regulating gene drives. *Science Policy Forum*. Published Online July 17, 2014. <https://doi.org/10.1126/science.1254287>.
- Palmer, C. (2010). *Animal ethics in context*. New York: Columbia University Press.
- Parfit, D. (1976). On doing the best for our children. In M. D. Bayles (Ed.), *Ethics and population* (pp. 100–115). Cambridge, MA: Schenkman Publishing Company Inc.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Paul, L. A. (2014). *Transformative experience*. Oxford: Oxford University Press.
- Pearce, D. (2009). *Reprogramming predators*. Published Online. <https://www.hedweb.com/abolitionist-project/reprogramming-predators.html>. Last updated 2015. Accessed February 15, 2017.
- Pearce, D. (2016). *Compassionate biology. How CRISPR-based "gene drives" could cheaply, rapidly and sustainably reduce suffering throughout the living world*. Published Online. <https://www.hedweb.com/gene-drives/index.html>. Accessed February 15, 2017.
- Rachels, J. (1975). Active and passive euthanasia. *New England Journal of Medicine*, 292, 78–80.
- Regan, T. (2004). *The case for animal rights [1983], Second Edition updated with a new preface*. Berkeley: The University of California Press.
- Reznick, D., Bryant, M. J., & Bashey, F. (2002). *r*- and *K*-selection revisited: The role of population regulation in life-history evolution. *Ecology*, 83, 1509–1520.
- Rolston, H. (1988). *Environmental ethics: Duties to and values in the natural world*. Philadelphia: Temple University Press.

- Rolston, H. (1989). Duties to endangered species. *Philosophy gone wild* (pp. 206–222). Amherst, NY: Prometheus.
- Sandler, R. (2012). *The ethics of species*. Cambridge: Cambridge University Press.
- Sapontzis, S. F. (1987). *Morals, reason, and animals*. Philadelphia: Temple University Press.
- Savulescu, J., & Persson, I. (2012). *Unfit for the future: The need for moral enhancement*. Oxford: Oxford University Press.
- Simberloff, D. (1998). Flagships, umbrellas and keystones: Is a single-species management passé in the landscape era? *Biological Conservation*, 83, 246–257.
- Simmons, A. (2009). Animals, predators, the right to life, and the duty to save lives. *Ethics & the Environment*, 14, 15–27.
- Singer, P. (1975). *Animal liberation*. New York: Avon Books.
- Slovic, P. (2007). ‘If I look at the mass I will never act’: Psychic numbing and genocide. *Judgment and Decision Making*, 2(2), 79–95.
- Sözmen, B.İ. (2013). Harm in the wild: Facing non-human suffering in nature. *Ethical Theory and Moral Practice*, 16(5), 1075–1088.
- Thomson, J. J. (1976). Killing, letting die, and the trolley problem. *The Monist*, 59(2), 204–217.
- Tomasik, B. (2015). The importance of wild animal suffering. *Relations*, 3, 133–152.
- Tomasik, B. (2017). Habitat loss, not preservation, generally reduces wild-animal suffering. Manuscript available at *Essays on Reducing Suffering*. <http://reducing-suffering.org/habitat-loss-not-preservation-generally-reduces-wild-animal-suffering/>. Accessed October 17, 2017.
- Vitt, L. J., & Caldwell, J. P. (2009). *Herpetology: An introductory biology of amphibians and reptiles* (3rd ed.). Burlington: Academic Press.
- Walker, B., Holling, C. S., Carpenter, S. R., & Kinzig, A. (2004). Resilience, adaptability and transformability in social-ecological systems. *Ecology and Society*, 9(2), 5.
- Walker, B., & Salt, D. A. (2006). *Resilience thinking: Sustaining ecosystems and people in a changing world*. Washington, D.C.: Island Press.
- Weisberg, M. (2006). Robustness analysis. *Philosophy of Science*, 73(5), 730–742.